CONFERENCE
PROCEEDINGS

Understanding the Insider Threat

Proceedings of a March 2004 Workshop

Richard C. Brackney, Robert H. Anderson

NATIONAL SECURITY RESEARCH DIVISION

CONFERENCE PROCEEDINGS

Understanding the Insider Threat

Proceedings of a March 2004 Workshop

Richard C. Brackney, Robert H. Anderson

Prepared for the Advanced Research and Development Activity

NATIONAL SECURITY RESEARCH DIVISION

The work described here was conducted in the RAND National Security Research Division, which conducts research and analysis for the Office of the Secretary of Defense, the Joint Staff, the Unified Commands, the defence agencies, the Department of the Navy, the U.S. intelligence community, allied foreign governments, and foundations. These proceedings were supported by the advanced information research area in the Advanced Research and Development Activity within the U.S. intelligence community.

ISBN 0-8330-3680-7

Published 2004 by the RAND Corporation
1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138
1200 South Hayes Street, Arlington, VA 22202-5050
201 North Craig Street, Suite 202, Pittsburgh, PA 15213-1516
RAND URL: http://www.rand.org/
To order RAND documents or to obtain additional information, contact
Distribution Services: Telephone: (310) 451-7002;
Fax: (310) 451-6915; Email: order@rand.org

Preface

The Advanced Research and Development Activity (ARDA) within the U.S. intelligence community (IC) has several research "thrusts," including one on advanced Information Assurance (IA) headed by Richard C. Brackney. On March 2–4, 2004, an unclassified workshop was held at the offices of McAfee Security (a division of Network Associates, Inc.) in Rockville, MD. The topic was "Understanding the Insider Threat."

The format of the workshop combined plenary sessions and four "breakout" groups, whose specialized topics were the following:

- Intelligence Community (IC) System Models
- Vulnerabilities and Exploits
- Attacker Models
- Event Characterization.

The workshop brought together members of the IC with specific knowledge of IC document management systems and IC business practices; persons with knowledge of insider attackers, both within and outside the IC; and researchers involved in developing technology to counter insider threats.

These proceedings contain an overview of the findings from this workshop and the display charts from briefings given to workshop participants. This document should be of interest to researchers investigating methods for countering the insider threat to sensitive information systems, and to members of the intelligence community concerned with the insider threat and its mitigation.

The RAND Corporation's research for ARDA's IA thrust is conducted within the Intelligence Policy Center (IPC) of the RAND National Security Research Division (NSRD). RAND NSRD conducts research and analysis for the Office of the Secretary of Defense, the Joint Staff, the Unified Commands, the defense agencies, the Department of the Navy, the U.S. intelligence community, allied foreign governments, and foundations.

For more information on the Intelligence Policy Center, contact the Acting Director, Greg Treverton. He can be reached by e-mail at Greg_Treverton@rand.org; by phone at (310) 393-0411; or by mail at RAND, 1776 Main Street, Santa Monica, CA, 90407-2138. More information about RAND is available at www.rand.org.

Contents

Figures

Figures

Tables

Summary

A major research thrust of the Advanced Research and Development Activity (ARDA) of the U.S. intelligence community (IC) involves information assurance (IA). Perhaps the greatest threat that IA activities within the IC must address is the "insider threat"—malevolent (or possibly inadvertent) actions by an already trusted person with access to sensitive information and information systems.

This unclassified workshop, held March 2–4, 2004, focused on the insider threat and possible indicators and warnings, observables, and actions to mitigate that threat. The ARDA researchers participating gave special attention to the activities, processes, and systems used within the intelligence community.

A combination of plenary and breakout sessions discussed various aspects of the problem, including IC system models, vulnerabilities and exploits, attacker models, and characterization of events associated with an insider attack. A set of presentations by members of the IC and its contractors on Intelink (Appendix G) and such research activities as the development of "Glass Box" software (see Appendix H) and ARDA's "Novel Intelligence from Massive Data" (NIMD) research program (Appendix I) aided the workshop discussions. The present workshop built upon the availability of materials generated in an earlier workshop focused on the insider threat (Appendix F).

Several overall themes emerged from these deliberations, discussed below under the headings of "Research Questions and Challenges" and "Databases Needed" (by researchers).

Intelligence Community System Models

The overall intelligence process involves requirements, collection, processing and exploitation, analysis and production, dissemination, and consumption, with feedback loops at all steps, as shown in Figure S.1.

Variant models, such as the NSA Reference Model (NRM), also exist. Of key concern to this group of researchers was the question: What "observables"[1] can be obtained at all stages of this process that would allow comparison of normal analyst activity with abnormal activity—which is potentially, but not necessarily, malevolent? Figure S.2 provides an indication of the richness of the concept of "observable"; it is a taxonomy developed by the earlier insider threat workshop cited above. Similar taxonomies characterize IC "assets" and "users."

[1] An *observable* is anything that can be detected with current technology. A number of workshop participants argued that this definition should be broadened to include foreseeable future technological developments.

Figure S.1
Intelligence Process

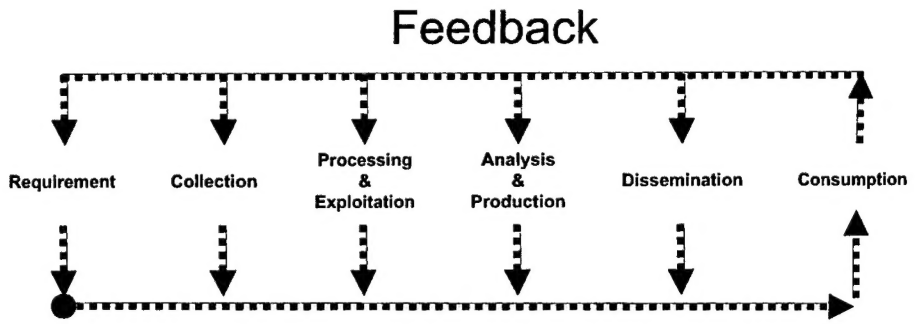

Figure S.2
Taxonomy of Observables

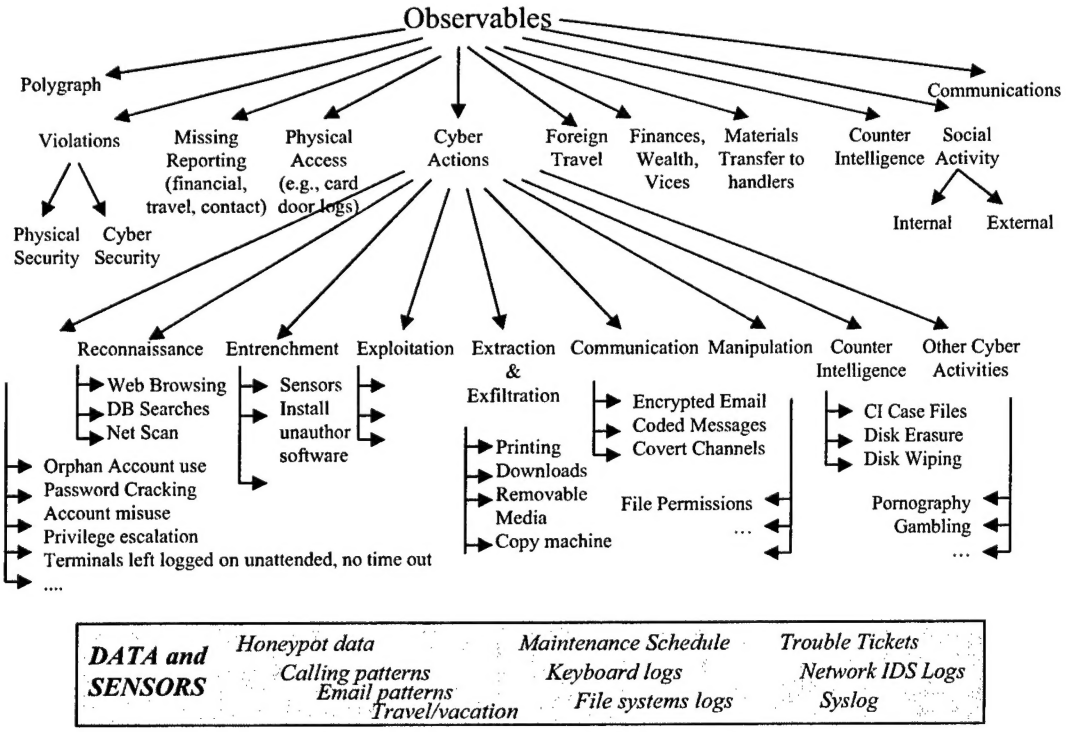

Vulnerabilities and Exploits

What types of exploits[2] might an insider use to obtain information, alter its integrity, or deny its availability to those who need it? This workshop concentrated on cyber-related

[2] The noun *exploit* is often used within the intelligence community to mean the development of a plan (and, usually, its subsequent execution—often surreptitiously) to obtain information or an advantage.

exploits because they were felt to be potentially the most damaging and most likely to increase in the future, as a new generation of analysts emerges with more computer skills than the previous generation.

Workshop participants generated a list of 33 example exploits. For each they listed a brief description, preconditions that would allow the exploit to happen, observables that might be generated during the exploit, and effects of the exploit (usually one of the following: a breach of confidentiality, integrity, or availability, or an enabler of other exploits). The short titles of the vulnerabilities are listed in Table S.1. Further details may be found in Chapter Three.

Attacker Models

Figure S.3 shows an overall model of the steps involved if a malevolent insider were to "mount an attack" against an IC asset. The attack might be as simple as obtaining access to information he or she does not have a need to know or as complex as disabling a key intelligence collection/processing/dissemination system.

Another way of depicting attacker actions is shown in Figure S.4. Here the attacker steps—motivation, benefit/risk assessment, acquiring the "client," collecting payment—were

Table S.1
Vulnerabilities and Exploits

1. Virus-laden CD and/or USB flash drive and/or floppy	18. Mislabeled paper
2. Administrator lockout	19. Netmeeting/WebEx controls
3. Social engineer passwords	20. "Day zero" attacks based on source code availability
4. Retry Internet attacks	21. Covert channels through steganography[a]
5. Smuggling out USB flash device or other media (exfiltration)	22. Copy and paste between classifications (from high to low)
6. "Missing" laptops/hardware	23. Internal e-mail that performs attacks
7. Targeted acquisition of surplus equipment	24. Wireless telephone cameras to capture information
8. Unpatched systems	25. Telephone tap recording onto removable media
9. Sabotaged patches	26. Telephone tap via hacking PBX telephone controller
10. False positives on anti-virus	27. Analyst changes workflow to exclude other analysts (dissemination)
11. Use of unattended terminal	28. Analyst changes workflow to include himself/herself
12. Targeting database "adjustments"	29. Insert bad content into report upon inception (e.g. translation)
13. Install software on host computer to capture keystrokes logger	30. Delete/withhold content into report upon inception
14. Extra copy of DB backups	31. Redirect analyst resources to support adversary's agenda
15. Wireless transmissions	32. Poor quality analysis/results/ reports
16. Cell phone/PDA/voice recorder in classified meeting	33. Get IC asset to collect info that benefits an unauthorized party
17. Suspicious activity on real systems (e.g., searching own name in databases)	

[a]Steganography is the hiding of information by embedding in an innocuous message or file, such as a digitized picture.

Figure S.3
Spiral Model Flowchart

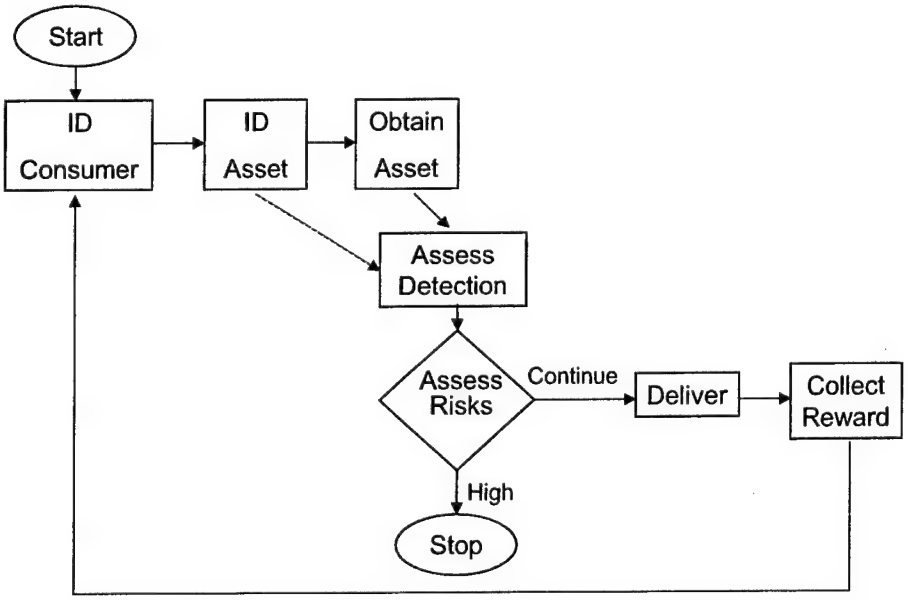

Figure S.4
Insider Attack Actions (white items not cyber observable)

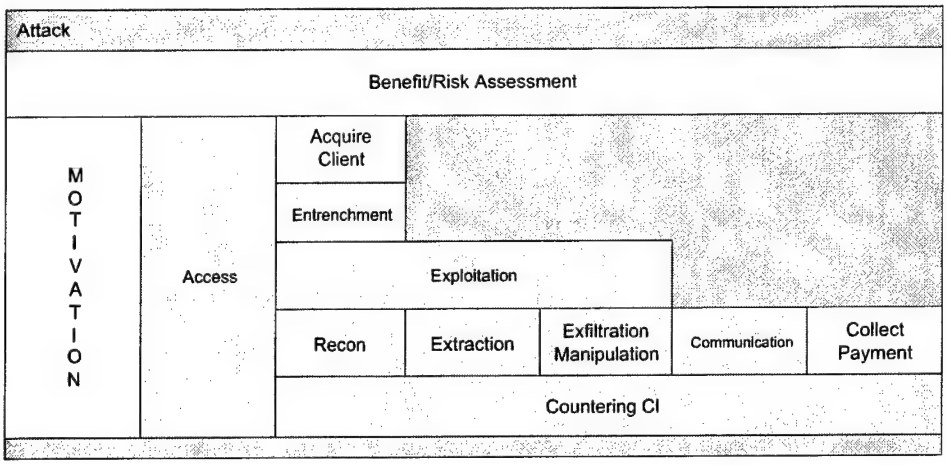

deemed *not* to generate cyber observables (that is, they would not be detected by information systems now in use or with enhancements planned by researchers and developers).

Given the various steps an attacker follows, as shown in Figure S.4, which steps are candidates for using the vulnerabilities and exploits shown in Table S.1? The answer is shown in Figure S.5, where the unitalicized insider actions have parenthesized numbers linking them to numbered entries in Table S.1. The parenthesized suffix letters C, I, A, E indicate whether the actions would lead to a breach of information Confidentiality, Integrity, Availability, or would be an Enabler of other attacks.

Figure S.5
Insider Actions Taxonomy Cross-Referenced with Vulnerabilities and Exploits (V&E) List

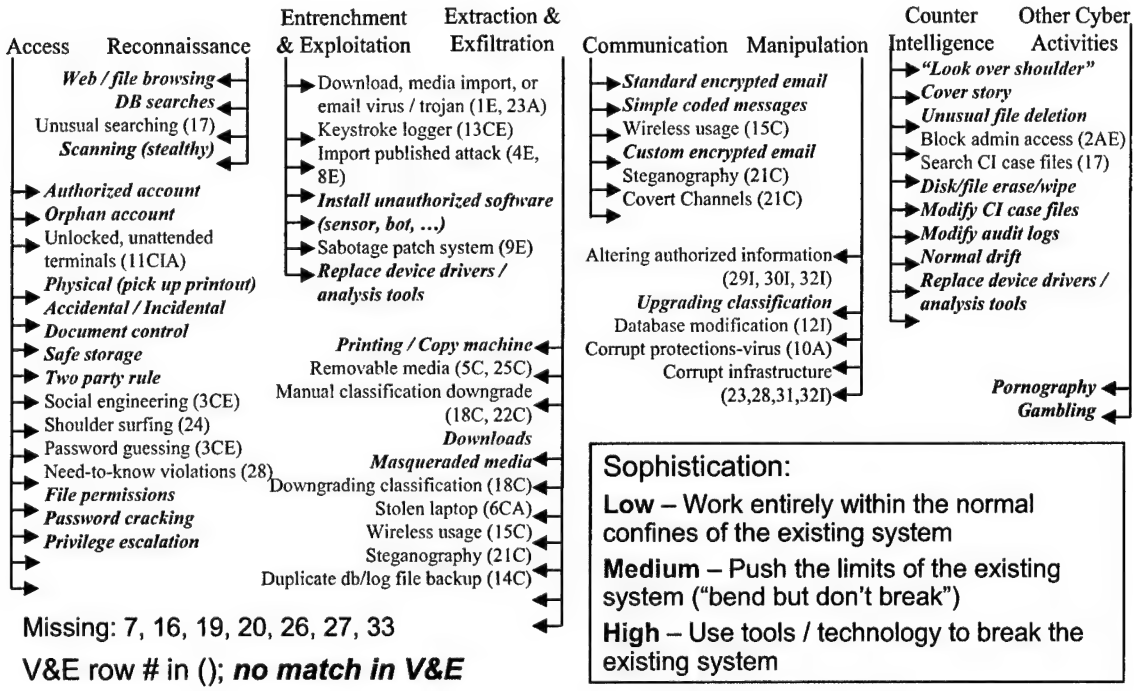

Event Characterization

As attacker actions generate observables through the operation of "detectors" of those observables, indicators of possible abnormal activity are generated. Those indicators can form a report; multiple reports can be fused into an "incident"; and multiple incidents then fused into a "case" of one or more incidents.[3] That process is shown graphically in Figure S.6.

Research Questions and Challenges

Each breakout group tried to formulate a set of research questions arising from its deliberations. Some groups stated these questions in the form of "grand challenges" to be addressed. We summarize the key questions and challenges below.

Six Categories of Research Questions

Research issues tended to fall within six categories:

1. User roles
2. Actions

[3] We assume that a *"case"* may be merely a collection of incidents having some commonality to be watched, or it could be the result of a post-facto analysis of source, cause, damage, etc.

Figure S.6
Data Collection Steps Regarding an Event

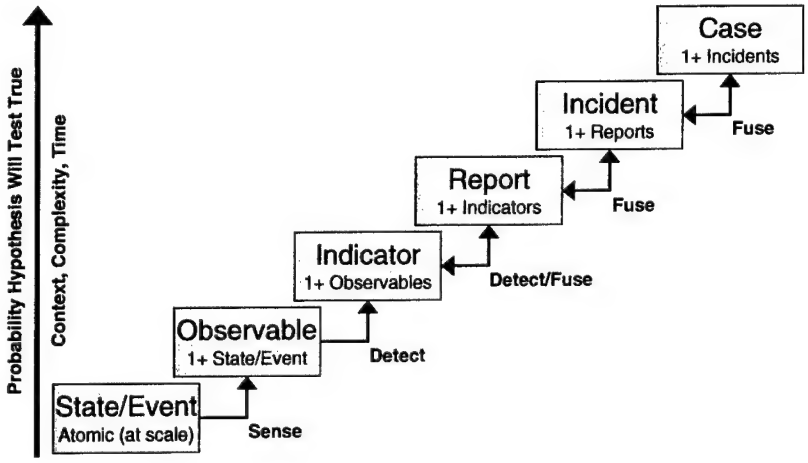

3. Observables (events)
4. Sensors
5. Fusion and analysis (both spatial and temporal)
6. "Triggers" (priorities, and level of certainty).

The first four categories each require *languages to describe them*, and *means for mapping each into the next* (i.e., from a description of user roles to a set of described user actions, which in turn lead to a set of potential observables. Those observables are then sensed and the sensed signals fed into fusion and analysis programs, which in turn create actions and alerts within the system).

An additional common thread is the need for *correlation and management tools* to correlate multiple events or triggers with an incident, to correlate multiple events with a case, and to correlate multiple cases into a coordinated attack.

The topic of sensors (item 4 in the above bulleted list) requires substantial research in at least the following areas:

- Identification of information that should go into an event record
- Development of sensors specific to particular applications
- Standardization of event record syntax and semantics; scales of severity and confidence; system interfaces; and means for establishing an inviolate "chain of evidence"
- Detection of "low and slow" attacks
- Optimization of selection, placement, and tuning of sensors
- Tradeoffs in adaptability: How do you recognize legitimate changes in user behavior? How do you resist the "conditioning" of sensors by a malicious insider (through a pattern of actions that "migrate" the sensor from a nominal setting to one that won't recognize the attack)?
- Development of validation and test data and techniques (see "Databases Needed," below).

Challenges

Participants stated several "grand challenges" for researchers:

- Define an *effective way of monitoring* what people do with their cyber access, to identify acts of cyber espionage. Focus on detection, not prevention. Such monitoring (or the perception of monitoring, which may suffice in some cases) can be an effective deterrent.
- Develop *policies and procedures* to create as bright a line as possible between allowed and disallowed behaviors (i.e., reduce the ambiguity).
- Consider *sociological and psychological factors* and create better cooperation between information systems personnel and human resources personnel (including security, medical, financial, and other support services). In short, broaden oversight of all aspects of a user's background and behaviors.
- *Combine events from one or more sensors* (possibly of various types or different levels of abstraction) to facilitate building systems that test hypotheses about malicious insider (MI) activity, to detect MI activity that is not detectable using a single event record, to develop a "calculus of evidence," to develop metrics for comparing and weighting diverse inputs, and to determine how "this fusion" can be used to create useful synthetic/compound events.

Databases Needed

Breakout sessions considered what databases would aid in this research if they were available. Researchers need databases containing examples of specific attacks, the characterization of normal behavior for users in different roles (including that of a system administrator), and artificial or real sensor data that include a mix of legitimate and malicious activity. Potential sources for the development of such datasets include a MITRE dataset of normal, and "insider threat" network activities; data from the ARDA NIMD[4] study; data obtained from use of the Glass Box[5] software; synthetically generated data from a simulator; and individual datasets developed by researchers that might be traded among projects.

A Concluding Remark

During a concluding plenary session, a senior member of the intelligence community, hearing the results from the various breakout session deliberations, made the comment, "What you're doing is important, but don't forget that IC analysts are people, too, and need a good work environment in which to stay motivated in their stressful jobs. When considering 'observables' and sensors and other means of keeping track of the activities of 'insiders,' please ask yourselves, 'Would I want to work in that (resulting) environment?'" It's important to keep this in mind, in the research enthusiasm for what *might* be monitored, and observed, and data-correlated. We must strike a balance between effectiveness in thwarting

[4] See Appendix I for information about the ARDA "Novel Intelligence from Massive Data" (NIMD) research thrust.

[5] See Appendix H for information about the "Glass Box" research effort.

insider exploits against intelligence assets and effectiveness in the process of generating and disseminating that intelligence information itself.

Acknowledgments

A three-day intensive workshop such as the one documented here requires substantial planning. The planning committee for this Insider Threat workshop consisted of Richard Brackney and John Farrell (ARDA), John C. Davis (Mitretek), Lisa Yanguas (NSA/R6), Paul Esposito (NSA/Defensive Computing Research Office), Tom Haigh (Adventium Labs), and Robert H. Anderson (RAND). Tom Haigh provided substantial help in organizing the summary of overall research issues and challenges emerging from the workshop.

Hosts for the workshop, providing excellent services and facilities, were Erik G. Mettala and David Sames (McAfee Research).

The organizers of the workshop also greatly appreciate the time and attention of senior members of the intelligence community who gave briefings on various aspects of the intelligence process and on research underway.

RAND colleague Diane C. Snyder provided very useful comments on a draft of this report.

Abbreviations

ACS	Automated Case System (FBI)
API	Application Program Interface
ARDA	Advanced Research and Development Activity
CD	compact disk
CD-ROM	compact disk–read-only memory
COI	community of interest
COMINT	communications intelligence
DCI	Director of Central Intelligence
DIA	Defense Intelligence Agency
DoS	denial of service
EUID	electronic user identification
HUMINT	human intelligence
H/W	hardware
IA	information assurance
IC	intelligence community
IR	infrared
LAN	local area network
MASINT	measurement and signatures intelligence
MI	malicious insider
NIMD	Novel Intelligence from Massive Data
NRM	NSA Reference Model
NSA	National Security Agency
NT	Windows NT (operating system)
OS	operating system
PBX	private branch exchange (telephone control)
PDA	personal digital assistant
PKI	public key infrastructure
QoS	quality of service

RF	radio frequency
RFID	radio frequency identification
RUID	radio frequency user identification
SAM	surface-to-air missile
SCIF	secure compartmented information facility
S/W	software
tcpdump	transmission control protocol dump (program)
TS/SI	top secret/special intelligence
URL	universal resource locator
USB	universal serial bus (computer port)
WMD	weapons of mass destruction

CHAPTER ONE

Introduction

The operations and analyses of the United States intelligence community (IC)[1] are based heavily on a set of information systems and networks containing extremely sensitive information. Most observers believe that the greatest threat to the integrity, confidentiality, and accessibility of the information in these systems is the "insider threat."[2] This phrase usually refers to a malicious insider, acting either alone or in concert with someone "on the outside" of these systems. However, one should also consider the possibility of unintentional actions by an insider that can have substantial adverse consequences or that draw attention to himself when innocent.

Discussions of the "insider threat" raise many questions: Who, exactly, is an insider? Anyone with physical or electronic access to these networks, including maintenance and custodial personnel? How much sophistication (if any) does it take to compromise the information within these systems? What defenses, including "indicators and warning," might be instituted to guard against this insider threat?

To address these questions, the Information Assurance (IA) research thrust of the IC's Advanced Research and Development Activity (ARDA) held a workshop on March 2–4, 2004. Participants included ARDA contractors working on the insider threat to information systems and members of the U.S. intelligence community with knowledge about its systems and networks. It was held at the offices of McAfee Security, a division of Network Associates, Inc., in Rockville, MD. The stated objectives of this workshop were:

---

[1] The agencies normally considered to constitute the IC are the office of the Director of Central Intelligence, the Community Management Staff, the National Intelligence Council, a set of Defense Agencies (Defense Intelligence Agency; National Security Agency; National Reconnaissance Office; Army Intelligence; Coast Guard Intelligence; Navy Intelligence; Air Force Intelligence; Marine Corp Intelligence; National Geospatial-Intelligence Agency—formerly the National Imagery and Mapping Agency), and the non-Defense agencies (Central Intelligence Agency; Federal Bureau of Investigation; Advanced Research and Development Activity; and portions of the Department of Treasury, Department of Energy, and Department of State.)

[2] As evidence for this statement, consider the following excerpt from a presentation on the Robert Hanssen case presented during the opening plenary session: (1) "Since the 1930s, every U.S. agency involved with national security has been penetrated by foreign agents, with the exception of the U.S. Coast Guard" (Webster Commission, 2002); (2) 117 American citizens have been prosecuted for espionage between 1945 and 1990 (or there is clear evidence of their guilt). Money appears to be the main factor; most spies volunteered their services. Prominent examples of insider spies include:

- *Aldrich Ames*, CIA counterintelligence officer (nine years as spy)
- *Ronald Pelton*, former intelligence analyst for NSA
- *Jonathan Pollard*, military intelligence analyst, gave Israel 800 classified documents, 1,000 cables
- *John Walker*, retired naval officer, with son and brother, supplied the Soviets with cryptographic material.

- To generate and capture domain knowledge that will benefit the broad base of researchers studying the Insider Threat. This includes, but is not limited to, knowledge about:
  - Inside attacker characteristics, including the vulnerabilities they tend to exploit, and the attack methods they use.
  - Attack characterization, including the necessary or likely preconditions for an attack, the observables generated during an attack, and the effects of the attack.
  - The electronic network and application systems used by the IC for document management, including the mechanisms used to protect the systems and data.
  - IC business models for generating and controlling access to documents.
- To foster cooperation among researchers by developing, to the extent it is practical, methods for describing common aspects of their work, such as event characterization, attack and attacker classification, etc.
- To focus researchers on specific systems and problems of interest to the IC. We expect these to take the form of challenge problems.

As can be seen from the above description, researchers investigating means to counter the insider threat formed the "target audience" for the workshop: its purpose was to supply them with relevant knowledge about the workings of the IC, the types of document or information processing used by IC analysts, and the architecture of the IC's underlying information networks.

The workshop was unclassified, requiring that only generic information about some aspects of IC information processing activities were transmitted to researchers. The intent throughout the planning for this workshop was that the information generated (and as captured in this present document) should be widely available to anyone working on the insider threat problem, without restrictions.

The remainder of this document consists of the results of the deliberations of the four breakout groups. (Those results were originally presented to the workshop on PowerPoint charts; they have been converted to a prose form for greater readability and uniformity of presentation in these proceedings.) The descriptions and charters given to those breakout groups were as follows:

- **IC Systems and Business Models** for generating and controlling access to documents. This group will capture core knowledge about the business processes and the supporting network and application systems used by the IC for document management (creation, update, and dissemination). This includes the physical, procedural, and technical mechanisms used to protect the systems, services, and information. Since the systems are highly heterogeneous, with different processes and mechanisms, depending on specific system functions, we expect this group will generate a family of system models reflecting current IC systems practices and anticipating future IC systems and practices.
- **Vulnerabilities and Exploits.** This group will collect and organize knowledge about the ways insiders have attacked systems in the past and the ways they might attack them in the future. The group will identify ways that insiders have exploited technical and procedural vulnerabilities in the past to compromise classified information or

to affect the integrity of critical information. The group will identify the necessary or likely preconditions for an attack, the observables generated prior to and during an attack, and the effects of the attack. It is important to emphasize that the charter of this group goes beyond studying and describing technical exploits that an attacker could use to "hack" the system. In the past many of the most damaging exploits have resulted from legitimate use of system accesses for illegitimate purposes.

- **Attacker Models.** This group will direct its attention to identifying and understanding the relevant behavioral characteristics of inside attackers. Examples of these characteristics are attacker objectives, level of system knowledge and access, level of patience, tolerance for detection risk or attack complexity, social engineering skills, and technical capabilities. The working group will not concern itself with underlying psychological, political, or economic factors that might motivate some of the attacker's behavioral traits. An important part of this group's effort will be to understand how the attacker's observable behavior can be used to identify him as an attacker.
- **Event Characterization.** This group will identify the key elements necessary to characterize events associated with insider attacks, to facilitate tracking and interpreting a potential insider attacker's activities. This will assist researchers who are trying to integrate input from a variety of sensors to assess the likelihood of attacker activity and likely attacker intent. It will also help sensor researchers know what capabilities to include in the sensors they define.

The appendices contain the invitation to the workshop, the agenda, a set of links to relevant "read-ahead" material, and a list of participants. We also include PowerPoint charts used in the following plenary presentations made by members of the intelligence community and their contractors:

- *The Robert Hanssen Case: An Example of the Insider Threat to Sensitive U.S. Information Systems,* by Robert H. Anderson, RAND Corporation
- An overview of the results of a recent ARDA workshop on Cyber Indications and Warning, by Mark Maybury, MITRE Corporation
- *Intelink Factoids,* by Peter Jobusch, Intelink Management Office
- *Glass Box Analysis Project,* by Frank L. Greitzer, Battelle, Pacific Northwest Division
- *Interacting with Information: Novel Intelligence from Massive Data (NIMD),* by Lucy Nowell, ARDA.

CHAPTER TWO

IC System Models

Relevant Taxonomies

This breakout group[1] began by reviewing a set of taxonomies developed in a previous ARDA "Indicators and Warning" workshop (see Appendix F), shown in Figures 2.1 through 2.3. These figures list a set of "observables" that might be used to determine abnormal behavior of an insider or of IC documents, "assets" within the IC that might be tracked, and a list of the different categories of "users" (insiders) within the IC community.

Figure 2.1
Observables Taxonomy

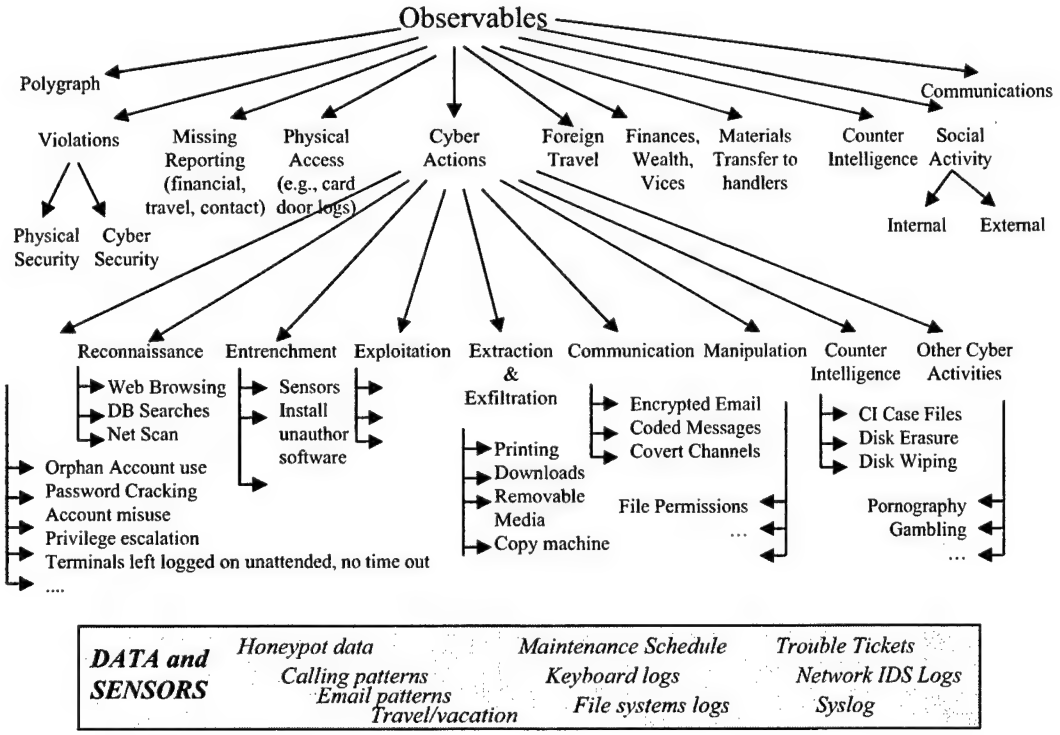

[1] Participants were Paul Esposito, Chris Geib, Joseph Giampapa, Alexander Gibson, Terrance (TJ) Goan, Clarence Jones, Jr., Linda (Miki) Kiyosaki, Sara Matzner, Mark Maybury, James Newton, David Sames, and Thomas Shackelford.

Figure 2.2
Assets Taxonomy

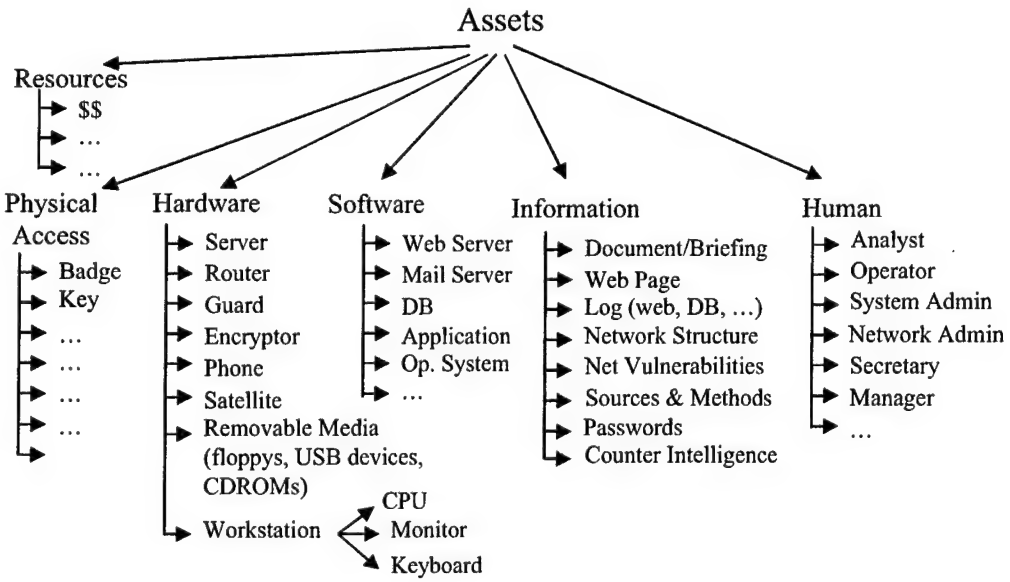

Figure 2.3
IC Users Taxonomy

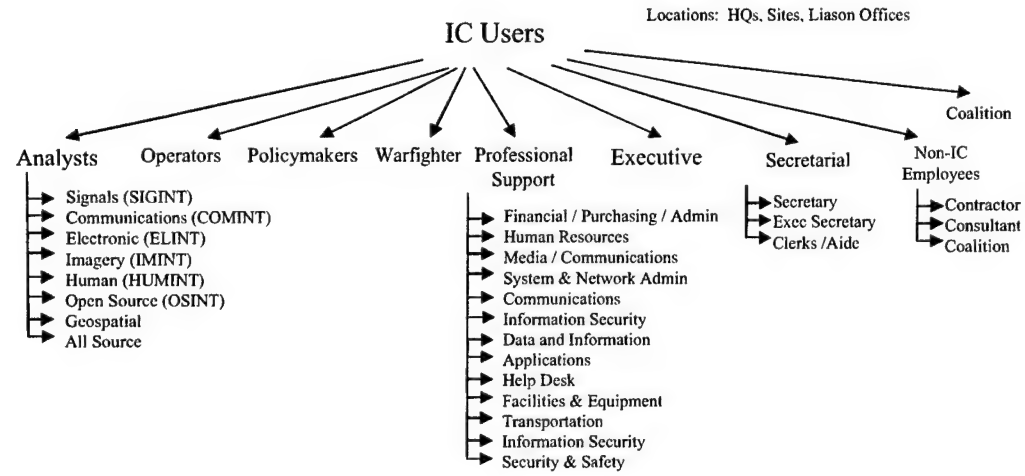

Focus: Discipline focus, Geospatial area, Targets
(e.g., people, organizations, locations)

The remainder of the group's deliberations then concentrated on a description of the intelligence process as it relates to a document life cycle, and a reminder that there are other systems involved to be considered: policy, personnel, physical security, etc.

Definition of the Term "Document"

The group developed the following definition of *"document,"* to be used in describing IC system and process models:

- Any collected artifact that is used to convey information.
- Ultimate purpose is to inform decisions at various levels:
  - — Strategic
    - Military
    - Legislative
    - Political
  - — Tactical
- Can be electronic or physical
- Can be structured or unstructured
- Image, voice, text, other
- Attributes
  - — Owned, managed, protected.

Characterization of the Intelligence Process

The group characterized the intelligence production process in terms of the diagram in Figure 2.4, involving requirements generation, collection, processing and exploitation, analysis and production, dissemination, and consumption, with various feedback loops.

In describing the terminology in Figure 2.4, the group also referred to the "NSA Reference Model" (NRM), which comprises the following steps:

- Signal, data information, knowledge, intelligence
- *Collection* gathers signals and data
- *Processing and exploitation* transform signals and data into information
- *Analysis and production* transform information into knowledge and intelligence
- Intelligence is *consumed*.

The following subsections describe each of these processes in more detail.

Figure 2.4
Intelligence Process

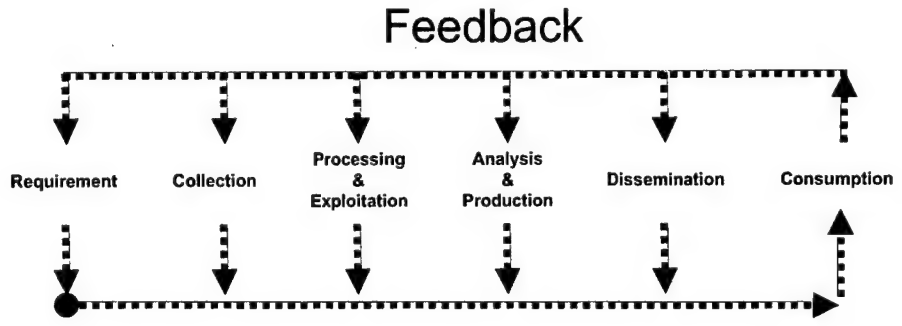

Requirement

A requirement is a statement of need by a consumer in the form of a formalized request. There are two types:

- *Standing requirement.* It is vetted by a consensus process, formalized by a memo, entered into a database, and is persistent.
- *Ad hoc requirement.* Anyone can submit an *ad hoc* requirement to a specific agency. It is stored in a database.

Some additional attributes of requirements are that they involve *checks and balances,* with multiple levels of vetting (e.g., committee meetings of analysts); there are *potential vulnerabilities* (e.g., the ability to change and modify requirements); the *internal threat level* is high; cyber or non-cyber attacks are possible on the database of requirements; and *indications and warning* of attacks on requirements could be derived from database audit logs.

Collection

The collection process was described as acquisition of raw data, which may include construction of new systems for performing the acquisition.

Its major components are all the "INTs" (e.g., COMINT, HUMINT, MASINT).

Additional attributes of the collection process include the existence of *checks and balances,* such as attribution of the source, techniques to preserve the integrity of collected data, and legal restrictions derived from government policy; *potential vulnerabilities*, such as degradation of collected data (including integrity issues), intercepts and eavesdropping, and denial of service from collecting and sending back what was collected; the *insider threat level* is considered to be "less likely"; and the *potential for collusion* was deemed to be "none."

Processing and Exploitation

The group defined processing and exploitation as selecting, filtering, and rendering the collected information into a human-usable form. It involves converting technically collected information into intelligence. Filtering is often involved and may be performed by individuals, software, or groups.

Analysis and Production

All IC member organizations perform analysis. Analysis is the transformation of information to knowledge. Production is the formalization of knowledge into a document or product.

The major *components* of analysis and production were listed as

- office automation tools
- secure document management systems
- specific analysis tools
- working aids, such as software search, visualization, and other programs
- communities of interest (COIs)
- other networked and local databases.

The process involves *checks and balances,* including hierarchical reviews for accuracy, consistency, accountability, attribution, security (e.g., assuring that proper markings are pre-

sent); collaborative production (not necessarily co-located); a coordination process involving interpretation and enforcement of policies; and authorizations.

The group felt that the analysis and production process had *potential vulnerabilities*, mainly to working documents and software programs constituting the working aids.

Dissemination

Dissemination is the distribution of intelligence to the requesters and authorized consumers. The group concentrated on electronic dissemination methods, with the following characteristics.

Its major components are chat, e-mail, and communities of interest. It uses both synchronous and asynchronous notifications and partial dissemination.

The group listed attributes of the dissemination process as

- *checks and balances* for dissemination, which are well-defined for paper-based documents and may use electronic watermarks for trace-back
- *potential vulnerabilities* including breach of confidentiality, denial of access to distribution lists, and distribution beyond intended consumers
- both the *internal threat level* and *collusion level* were considered to be high
- both cyber and non-cyber means of dissemination were considered to have *vulnerabilities*; cyber dissemination techniques were deemed to be less uniform
- *indications and warning* are to be considered on a per-dissemination channel basis.

Consumption

Consumption of intelligence is the use of produced intelligence by authorized users to support decisionmaking.

The main components of consumption (on the individual level) were stated by the group to be the five human senses.

Checks and balances in consumption include traditional security procedures and enforcement of policy. *Potential vulnerabilities* lurk during consumption, with the possibility of exfiltration, leaks, misuse, misinterpretation, and withholding.

An important research topic exists within the consumption process: finding ways to inject checks and balances within that process to provide observables.

The *internal threat level* related to consumption was deemed high; the *collusion potential* was deemed medium.

Although much dissemination is electronically based, consumption tends to remain primarily non-cyber, using low tech and traditional means (although the role of chat and e-mail is growing).

One *indicator and warning* of misuse of the consumption process is that restricted-dissemination data show up in the public press.

Definitions

The discussion group used the following definitions in describing the intelligence process:

- *Insider*: Anyone with access, privilege, or knowledge of information systems and services[2]
- *Malicious insider (MI)*: Motivated to intentionally adversely impact an organization's mission (e.g., deny, damage, degrade, destroy)
- *Observable*: Anything that can be detected with current technology[3]
- *Sensor*: Measures an observable (e.g., login, print, delete)
- *Sensor logs*: Recording of observables
- *Sensor stream*: Series of observables from a sensor
- *Indicator*: Identifiable event based on sensor output logs
- *Detect*: Determines an event based on processing of indicators
- *Report*: Indications and warnings of malicious insider behavior
- *Incident*: Related set of events
- *Fusion*: Processing multiple sensor outputs to provide an enhanced result (e.g., more abstract or more concrete; higher confidence)
- *Case*: One or more incidents that share common attributes, and are deemed to be (potentially) related.

Reference

The group cited Lowenthal (2003) as a reference for information about the intelligence process.

---

[2] Note that we do not say "legitimate access." Someone (e.g., a janitor, a service technician) may be given access accidentally or inadvertently but nevertheless have access to certain "insider" privileges.

[3] Since we are concerned with research on automated detection of insider threats, we do not include here observables that are only human-detectable, unless that observer acts like a sensor and records the observation for subsequent processing steps.

CHAPTER THREE

Vulnerabilities and Exploits

Group Focus

This breakout group[1] stated that their focus was on ways insiders attack information systems, including

- Preconditions for attack
- Observables
- Effects.

The group included illegitimate use of authorized access and focused on threats and vulnerabilities to IC networks. Because this was an unclassified workshop, certain vulnerabilities may be known at a classified level that cannot be described here; they will be described only at a generic level.

Overview of Group Deliberations

The group decided its limited time was best spent on the following activities:

- Look at some real life "war stories" about insider threats to critical information systems
- *Decompose* those and similar events to determine
  — *preconditions* (involving both physical and logical access)
  — *observables* (that could have been used to thwart the attack)
  — *effects* (of the attack).

"War Stories"

The group started by asking, "Has this ever happened in your world?" and gave these as examples. (The contributors vowed that all have been seen "in practice" in the real world, except for the second one, which was used in a test only.)

[1] Participants were Robert Anderson, Philip Burns, Matthew Downey, Jeremy Epstein, Dana Foat, Steve Harp, Dennis Heimbigner, Kevin Killourhy, Vincent Lee, Mark Morrison, Mike Pelican, and Brad Wood.

- An insider walks into the secure compartmented information facility (SCIF), pulls out from under his coat a freshly burned CD, sticks it into his *classified workstation*, and then selects RUN.
- An insider locks administrators out by making multiple attempts against their passwords until the system locks out their passwords.
- A user calls the help desk: "Hello, this is Major Smith—can you reset my password?" And the help desk doesn't verify that it's really Major Smith who's calling.
- An insider finds a nifty attack on the Internet. He asks himself, "Gee, I wonder if this will work on our LAN?"
- The business portion of the agency was sent a system patch, but they didn't give it to the security guys. The system didn't get patched.
- An insider modifies a valid system patch, which then gets distributed to the whole "world" of that agency via *LiveUpdate*.
- What if . . . someone modified a planning database to change the coordinates for SAM sites 2.0 km to the south, and make them SAM-2 instead of SAM-5 missiles (so that they were perceived to have shorter range). Pilots would get shot down. All this requires is access to a database in Microsoft Access, Excel, etc.
- A malicious insider copies a TS/SI file from his classified workstation onto a USB port "flash drive,"[2] moves it over to his unclassified system, and mails it out, all within the same office.
- Another insider installs a keystroke logger to get a few passwords to another computer in the same office.
- A database administrator makes an extra copy of the database files, but says the tapes are bad. He/she then carries the tapes out, and no one is the wiser.
- An insider has a wireless transceiver in his unclassified system, to transmit files after they have been moved from his classified workstation to his unclassified one (see "USB flash drive," above).

Attack Actions, Observables, Effects

Having "warmed up" on the above examples, the group then attempted to develop a more complete listing of "discrete attack actions" (many of which could be combined in various ways into more complete attack scenarios). In creating this list, the group used the following definitions:

- *Attack action*: Any nefarious activity undertaken by an adversary. (It does not have to result in a loss of confidentiality, integrity, or availability.) The group also intends to focus on "atomic" attacks that would be part of a larger campaign.
- *Observable*: Anything that could be detected with current technology, or with any other technology that might be considered possible. (Note that this definition extends the definition cited by the "IC System Models" group [Chapter Two]).

[2] The reference is to a very small keychain device that plugs into a computer's USB port and acts like a removable disk. At this writing (October 2004) they are available in sizes ranging from at least 32 KB to four gigabytes.

Table 3.1 at the end of this chapter contains the group's expanded list of "discrete attack actions." The table gives a name for each of the 33 attack mini-scenarios listed, some scenario details, preconditions, expected observables, and likely effects.

Reviewing the attacks listed at the end of this chapter, the group made the following general observations:

- Many are just *enablers* for chains of attacks.
- *Access* (either logical or physical) is a prerequisite condition for all attacks.
- Some attacks have *no observables*. It's an important research question to consider how that can be fixed.

In discussing the list of 33 attacks in a plenary session, the recommendation was made by a participant that this list should have an added column: "Existing remediation." That column would contain information on what measures are in place today, in various IC enclaves, to thwart the attacks listed. We recommend this as a useful piece of additional research to be performed. There was insufficient time for this group to investigate that issue and add the column during the workshop itself.

Another group used the listing in Table 3.1 and integrated these results into its own taxonomy (see Figure 4.11 in Chapter Four, "Attacker Models").

Roles

The group observed that the attacker's access and perspective vary depending on his role in the enterprise. Insiders could be system administrators, users, managers, analysts, linguists, "geeks" (computer specialists), or others.

In response to comments from the larger workshop that the attacks look technical, this group responded, *"They're [the attacks are] actually brain-dead!"* The group emphasized that although the attacks may look sophisticated to a "traditional" analyst, the next generation of analysts will have grown up with computers, cyber games, and the like, and all this will be second-nature to them. We need to think about *future* malicious insiders and not be overly influenced by previous attacks.

Grand Challenges

Each group attempted to formulate a set of "grand challenges" for research in discovering and mitigating the insider threat. The challenges listed by the "threats and vulnerabilities" group were the following:

- Create *effective deterrents* to cyber espionage.
  - —We need better ways to enforce and monitor the deterrents to put the "fear of God" in cleared people.
  - —The IC may not even need real monitoring; just the perception of monitoring may be enough (similarly, some states use cardboard state troopers to slow down traffic).

- Define an effective way of monitoring what people do with their cyber access, for purposes of identifying acts of cyber espionage.
  - —Focus on detection, not prevention (in the post-9/11 world, we need to allow everyone access to "everything"; instead, develop filters to find the nefarious acts).
- We need policies and procedures to create as bright a line as possible between allowed and disallowed behaviors to reduce the ambiguity.
  - —If the rules aren't realistic, then they dilute the overall impression of enforcement.
  - —The IC should therefore adjust the rules to be realistic and focus on what is important instead of trying to stop all disallowed behaviors equally.
- Consider *sociological/psychological factors,* and create better cooperation between information systems personnel and human resources personnel (to include security, medical, financial, and other support services). In short, broaden oversight of all aspects of a user's background and behaviors.
  - —Identify precursors to changes in an insider's "moral compass"—can this be modeled? Focus limited resources on insiders who present a greater risk.
  - —The clearance personnel should tell the cyber personnel who the risky people are (or what risky behavior is), and vice versa.
  - —We need multidisciplinary research teams (not just *geeks*) investigating what we should look for as indicators of possibly malevolent behavior.

Surprising Lessons Learned

Each group was asked, "What are the most surprising findings that came out of your deliberations?" This group answered:

- *Espionage case history does not cover cyberspace.* Most case histories do not involve interesting cyber exploits that we know could be used. Looking backward at case histories doesn't prepare us for what is coming with a more computer-savvy generation of analysts.
- *Things are looser than we might have expected.*
  - —Life is not as structured on the IC networks as we thought.
  - —Policy and practice aren't always the same (fewer people are searched, even sporadically and randomly, than expected).
  - —Our insiders are really trusted.

There is no practical way to prevent exfiltration by even a moderately determined adversary, especially given modern technology. For example, USB flash drives and CD-ROMs can hold huge amounts of data in a small space that can be hidden.

- COTS software is a real threat developed off-shore by uncleared foreign nationals.

Datasets Required

Another question asked of all groups was, "What datasets do you need for your research?" This group's responses were:

- The MITRE dataset might be extended to be more useful.
  - — Example: Record more things at the host level and more things outside the cyber domain beyond badge logs (e.g., where people are, use of photocopiers, phone records).
  - — "You don't need just one dataset—we need lots of them."
- Data from ARDA NIMD study[3] (recording "normal" analyst activity) seem very promising.
- Enhance the NIMD study with a different fictitious set of insiders (e.g., using some of the attack actions listed in Table 3.1) with other areas of interest and roles.
  - — Maybe leverage Glass Box software.[4]

Measures for Success

During their deliberations, this group asked themselves, "What are the measures by which we can judge success?" That is, how can we know that we have been successful at the end of the workshop? They decided on two criteria:

- We have identified observables that have not yet been highlighted by researchers
- We have developed a list of "challenge problems" based on the real threat to IC information systems.

Based on those criteria and the material in this chapter, the group felt its deliberations had been successful.

Table 3.1
Attack Actions, Preconditions, Observables, and Effects

Attack Action	Scenario Details	Preconditions	Observables	Effects
1. Virus-laden CD and/or USB flash drive and/or floppy	Malicious insider (MI) puts viruses or other malware on removable media, carries it into an IC environment, and inserts it into a system. Depending on the malware being introduced, it may impact confidentiality, integrity, and/or availability.	No physical checks on inbound materials; physical and logical access to machine; media on machine	Physical observation of media movement; tamper tape over media slots; closed-circuit camera, NT event log will show media access (but big impact on performance)	Enabler for numerous other attacks
2. Administrator lockout	MI finds names of administrators, and then tries to log in as the administrator (knowing that after some number of failures the admin will get locked out). Once all administrator accounts are locked out, the MI can perform attacks knowing that the admin is unlikely to be able to log in and detect or solve the problem.	Login access (remote or local), no multi-factor authentication, machine set to lock out after failed login attempts; names of admin user logins	Log entry for account lockout; log entry for admin lockout; repeated lockouts of any user over a short period of time	Enabler for numerous other attacks plus availability attack on SysAdmins

[3] See Appendix I for information about the ARDA "Novel Intelligence from Massive Data" (NIMD) research thrust.

[4] See Appendix H for information about the "Glass Box" research effort.

Table 3.1—Continued

Attack Action	Scenario Details	Preconditions	Observables	Effects
3. Social engineer passwords	MI calls up the help desk and says "Hi, I'm Major Smith, I forgot my password." Tries to convince help desk to reset or tell him/her the password.	Help desk doesn't have a way to absolutely authenticate request	Trouble ticket monitoring; integration with badging system to detect whether in building	Enabler for numerous other attacks plus confidentiality attack against the password itself
4. Retry Internet attacks	Attack scripts, worms, viruses, etc. from various websites in cyberspace are downloaded and executed on local secure LAN	Patches not installed promptly; if attack is trivial, network access, or if not, same as virus-laden CD or USB flash drive	Alien software installed on hosts; intrusion detection systems that monitor commands, system calls, URLs, etc.; frequency of patch installation; patches don't fix problems (hard to tell if up to date)	Enabler
5. Smuggling out flash drive or other media (exfiltration)	USB flash drive on keychain, or hidden on body during ingress and egress to controlled area	Insufficient physical checks on outbound; physical and logical access to machine; media on machine	NT event log shows media access; physical check on egress	Confidentiality
6. "Missing" laptops/ hardware	A laptop, personal digital assistant (PDA) or other device is removed from a secure facility; the MI can gain access to its contents	Insufficiently protected data on machine; physical access; aperiodic inventory checks	RFID on hardware devices (but also helps adversaries); failed physical inventory check; network census failure	Confidentiality + Availability
7. Targeted acquisition of surplus equipment	Bad guys buy surplus equipment from government agencies at auctions (perhaps tipped off by insider), and search disks for sensitive information	Insufficient sanitization process; physical access to equipment	"For sale" ad on eBay... photo with classified sticker	Confidentiality
8. Unpatched systems	An insider takes advantage of knowledge that sensitive info systems aren't patched promptly, and uses a recent attack method to gain root access to a server	Patches available but not installed promptly	Time interval between patching of operating systems, applications, etc.; vulnerability checkers	Enabler
9. Sabotaged patches	MI alters a patch to be disseminated to all LAN systems in the secure facility, enabling a trapdoor to permit greater access; that patch then gets installed automatically on all systems within the enclave	Ability to alter patches; patch distribution system	Integrity check with vendor to ensure patches are unchanged	Enabler
10. False positives on antivirus	MI creates a file containing the signature of a known virus, and distributes it within the enclave. Virus detection software sends alerts and restricts access, causing denial of service	Ability to create a virus signature (DoS attack)	Abnormally high number of virus infection indicators on atypical files	DoS/Availability
11. Use of unattended terminal	While a user is absent from his terminal, MI accesses his/her files, which are more sensitive and restricted than those to which he has normal access	Physical access; no timely screen saver/lockout; insufficient authentication	Abnormal signature from terminal	Confidentiality/ Integrity/ Availability
12. Targeting database "adjustments"	MI makes subtle changes to data within a key (e.g., targeting) database, causing inaccuracy in subsequent field operations	Database access; knowledge of database layout/semantics; no integrity checks	Transaction log indicating change	Integrity

Table 3.1—Continued

Attack Action	Scenario Details	Preconditions	Observables	Effects
13. Keystroke logger	MI installs a keystroke logger hardware or software on an associate's computer within the enclave; later retrieves password information from it	Physical or logical access (depending on type of logger); may require unmonitored access	Physical tamper indicator for physical logger; Tripwire-type software	Confidentiality + Enabler
14. Extra copy of DB backups	MI writes a backup copy of a sensitive database to tape; declares the backup or tape "bad," and reperforms backup. Removes extra copy of backup tape from the facility	Ability to make extra backup, physical access; no monitored egress; no accountability on media	Backup log (or change to backup log), # of tapes used, destruction evidence (not logs)	Confidentiality
15. Wireless transmissions	MI uses a wireless PC card, USB device, or other device to send wireless signals out of a sensitive enclosure, to an associate or device outside the facility	Ability to install or enable wireless; small wireless transceiver; enclosure isn't sufficiently shielded (IR, 802.11, other?)	Picking up wireless signals (RF), some optical techniques to detect IR (monitoring IR in an enclave might be enhanced by adding reflectors)	Confidentiality
16. Cell phone/PDA/voice recorder in classified meeting	Modern cellphones and PDAs are sufficiently small that they can be brought surreptitiously (or even accidentally) into a secure facility unnoticed; they can be put into "record" mode to capture voice recordings (e.g., into voice message files)	Insufficient physical monitoring; insufficient shielding; insufficient RF monitoring	Metal/device detector in/out of SCIF; RF detector in facility	Confidentiality
17. Suspicious activity on real systems (e.g., searching own name in databases)	Robert Hanssen periodically queried the FBI info systems for his own name, plus "dropbox" and other keywords, to see if he was being discovered	No monitoring of activity (capture and analysis of audit trails)	"Adversary activity" detected in audit logs	May implicate adversary
18. Mislabeled paper	A sensitive document has its header and foot labels changed (e.g., by using a word processor) to indicate a lesser (or no) degree of classification or sensitivity, is then printed	Possible to change markings on page; reliance on page markings as authorization for removal	No cyber observable! (maybe some cyber observable if there are highly instrumented applications/OS-level)	Confidentiality
19. Netmeeting/WebEx controls	Software such as "Netmeeting," "WebEx," or "Timbuktu" allows one user's display to be linked to, and slaved onto, another's. Can this be done surreptitiously without the first person's knowledge (e.g., when they're away from their terminal)? If so, all documents, e-mail, etc. displayed on the first terminal are mirrored onto the second one	Existence and default enablement of tools within enclave	Test if tool installed/enabled; test if tool active; indicate who is using tool (attribution). Determine what traffic is being exchanged using that tool	Enabler
20. Zero day attacks based on source	Major portions of source code for an earlier version of the Microsoft Windows operating system have become available on the Web. An MI uses knowledge gained from the source to craft a new attack, and uses it within the secure enclave. It is called "zero day" because it has never been used before, and hence no monitors are configured to detect it	Disclosure/knowledge of source code; tools to search for vulnerabilities; assumption that obscurity is significant source of security	No direct cyber observables; some derivative observables based on undesirable applications behavior; knowledge that source code has been disclosed	Enabler

Table 3.1—continued

Attack Action	Scenario Details	Preconditions	Observables	Effects
21. Covert channels through steganography	An MI hides sensitive information within a picture or other digital unclassified document, and removes it from the secure facility	Technology to allow encoding and decoding (including ability to install software)	Information content analysis results	Confidentiality
22. Copy and paste between classifications (from high to low)	An MI uses "copy "and "paste" between two windows on his/her PC, allowing one or more paragraphs from a highly sensitive document to be placed in a less-sensitive one, then removes the latter document	Ability to copy and paste between classification levels	Applications/OS instrumentation only (difficult problem)	Confidentiality
23. Internal e-mail that performs attacks	MI creates and sends e-mail within the secure facility that has JavaScript or an attachment with content that installs trapdoors or creates other vulnerabilities	E-mail agent processes scripting without user approval or knowledge	E-mail download with content; executable content; indication of whether the content/execution parsing agree; indication of whether the parsed content is nefarious (hard problem)	Enabler
24. Wireless cameras to shoulder surf	Installs a small, wireless videocamera in the ceiling above another user's workstation to capture keystrokes and obtain passwords	Ability to install camera and exfiltration (via wireless or media)	RF emissions	Confidentiality
25. Telephone tap recording onto removable media	A telephone tap is installed within a secure facility that records information on a CompactFlash card or other removable medium	Plain-text voice transmission; access to telco media	None	Confidentiality
26. Telephone tap via hacking PBX	Modern PBXs are computers. They can be "hacked" via default maintenance passwords or other means, to gain access to telephone conversations	Plain-text voice transmission; access to telco media; physical or logical access to PBX/distribution	Depends	Confidentiality
27. Analyst changes workflow to exclude other analysts (dissemination)	MI surreptitiously changes dissemination instructions (e.g., in software) to exclude some analysts from obtaining information they need, thereby harming the resulting intelligence product	Analyst ability to change/determine workflow for intelligence product (management authority)	Actual product workflow	Integrity/Availability
28. Analyst changes workflow to include himself/herself	MI surreptitiously adds him/herself to dissemination instructions to become aware of sensitive information without need to know	Analyst ability to change/determine workflow for intelligence product (management authority)	Actual product workflow	Integrity
29. Insert bad content into report upon inception (e.g., translation)	MI involved with the creation of intelligence information (e.g., during a translation process) alters its content, thereby harming the intelligence product	Translation authority (ability to create source report)	Inspection of product; comparison of product to raw data	Integrity
30. Delete/withhold content into report upon inception	MI involved with the creation of intelligence information (e.g., during a translation process) deletes key portions of its content, thereby harming the resulting intelligence product	Translation authority (ability to create source report)	None (probably); comparison to redundant report	Integrity

Table 3.1—continued

Attack Action	Scenario details	Preconditions	Observables	Effects
31. Redirect analyst resources to support some agenda	MI obtains access to software or other procedures by which analyst resources are allocated, and redirects them for own purposes	Analyst ability to change/determine workflow for intelligence product (management authority)	Collection does not meet exact requirement	Integrity/ Availability
32. Poor quality analysis/ results/ reports	MI involved in intelligence analysis deliberately creates faulty or misleading reports	Translation authority (ability to create source report)	None/few; customer feedback; job performance reviews	Integrity
33. Get IC asset to collect information that benefits an unauthorized party	MI obtains access to the process by which IC assets are tasked, and alters the tasking so that information of use to the MI (but not the U.S. intelligence effort) is obtained	Tasking authority	Lack of requirement for the collection activity.	Availability

CHAPTER FOUR

Attacker Models

Group Focus

This breakout group[1] began its deliberations by stating some broader concerns:

- Are attacker models providing too much information to the bad guys on what we look for?[2]
- Can individual-level activity monitoring cause the malicious insider to change but not limit his activity while causing the benign insider to restrict his creativity?
- How should we reconcile the need to collaborate more with better security?

The group then listed the following more specific questions about modeling attackers as important items to discuss:

- What are the defining characteristics and reasonable values for attacker models?
- What are the *observables* (actions, artifacts) that an insider generates and how can they be used to determine benign or malicious behavior?
- How could an attacker "cover his trail" of observables?
- What is the overall scope of activities that a malicious insider might undertake and what role do environmental factors play?
- Can we create a taxonomy of security controls, and how can we correlate these controls with detected observables to tune and manage responses?
- How can we generate models of normal behavior as well as malicious behavior, and how can we generate data to test the models?
- How can we find signs of malicious insider behavior in unusual environments, such as high-performance machines or massive visualization data transport?
- How do insiders seek information and how can researchers use context to determine if the information seeking is benign or malicious?

---

[1] Participants were Matthew Downey, Tom Haigh, Steve Karty, Van Lepthien, Rich Neely, Greg Stephens, Frank Greitzer, Thomas Hetherington, Stephen Laird, Tom Longstaff, Marisa Reddy, and Edward Wright.

[2] This was one of the few instances during the workshop in which counterintelligence effects of insider threat research were explicitly highlighted.

A First-Cut Notional Insider Model

Figure 4.1 represents a top-level view of an insider attacker model used by the group.

The group then compared this model to the known Robert Hanssen case history (see Appendix E for details). The Hanssen case was summarized as shown in Figure 4.2.

The group then elaborated the attacker model into the "spiral model" shown in Figure 4.3.

Figure 4.1
Notional Insider Model

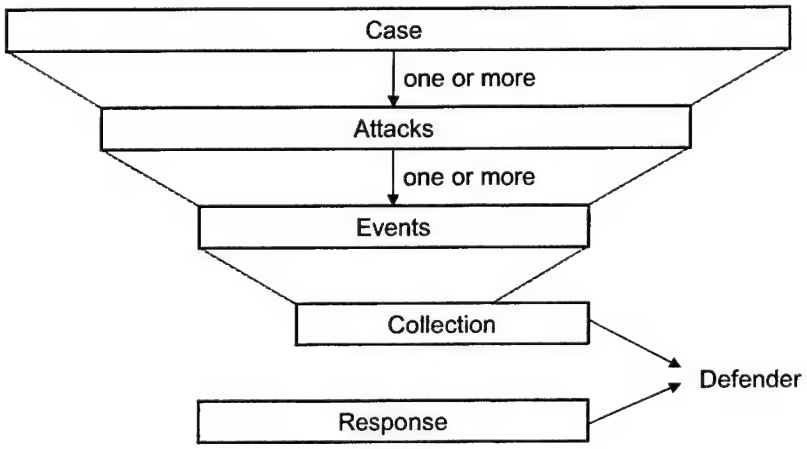

Figure 4.2
Hanssen Case History

1. Pre-violation – Granted Access
2. Trigger – Financial Status
3. Spiral Model (zero or more occurrences of each step, one or more events per attack)
    1. Identify Consumer / Recruitment / Motivation
        1. Search – Known (OJT)
        2. Contact – Approached
            1. Anonymous
            2. Attributed
        3. Negotiation – Handler / Instruction
    2. Identify Asset
    3. Look for Detection (assess risk)
    4. Obtain Asset
    5. Exfiltrate / Delivery
    6. Cover Tracks
    7. Collect Payment / Reward / Satisfaction

    } Attack
4. Exit

A user "case" was then abstracted into a set of actions occurring over time, as shown in Figure 4.4.

Figure 4.3
Spiral Model Flowchart

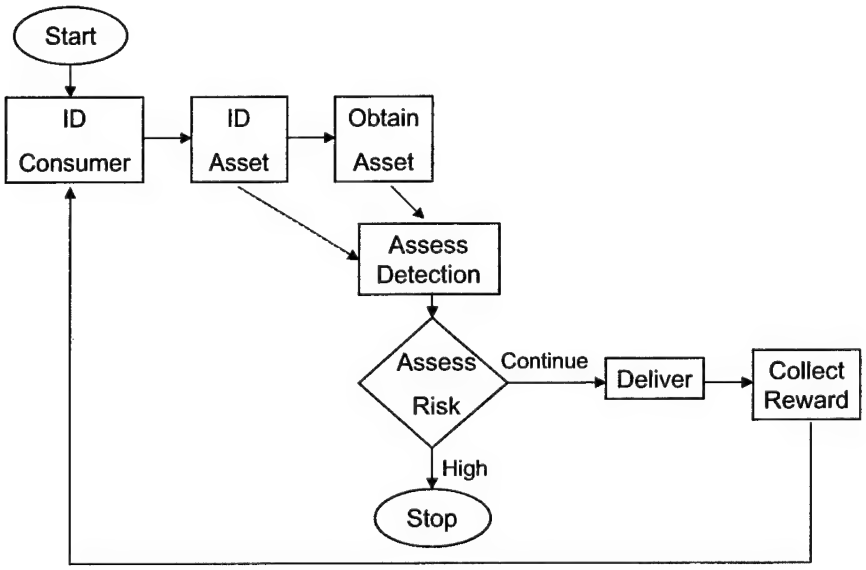

Figure 4.4
Insider Attack "Case" Actions Over Time

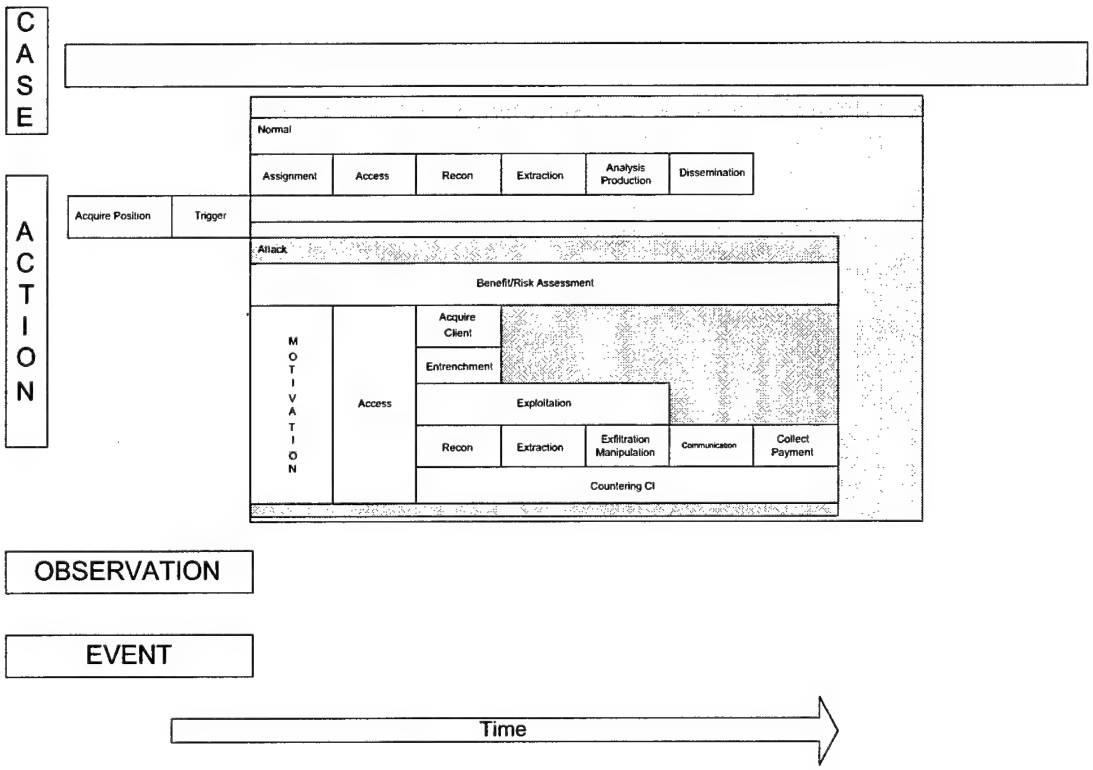

In contrast, normal user actions within an IC environment were diagrammed as in Figure 4.5.

The normal model (Figure 4.5) contains "base" behavior, and will continue while hostile activity is being done. It should be noted that normal behavior will be incorporated into most attacks. Normal behavior has analogues in hostile behavior—distinguishing the two can be a very hard problem.

Given the attack model in Figure 4.4, which of those actions generate "observables" that might be detected? Figure 4.6 shows that same model, with *unobservable* actions in white.

Another view of an attacker model is that the insider has certain attributes that are perhaps measurable. Attacks have certain observables, and the type of outcome can be placed in several categories. Figure 4.7 lists the attributes and categories itemized by this group's deliberations.

The group then created a modified taxonomy of insider actions, listing a number of possible actions that might be taken in the following categories:

- obtaining access
- reconnaissance
- entrenchment and exploitation
- extraction and exfiltration

Figure 4.5
Normal Insider Actions

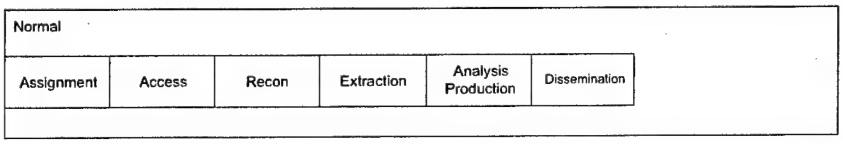

Figure 4.6
Insider Attack Actions

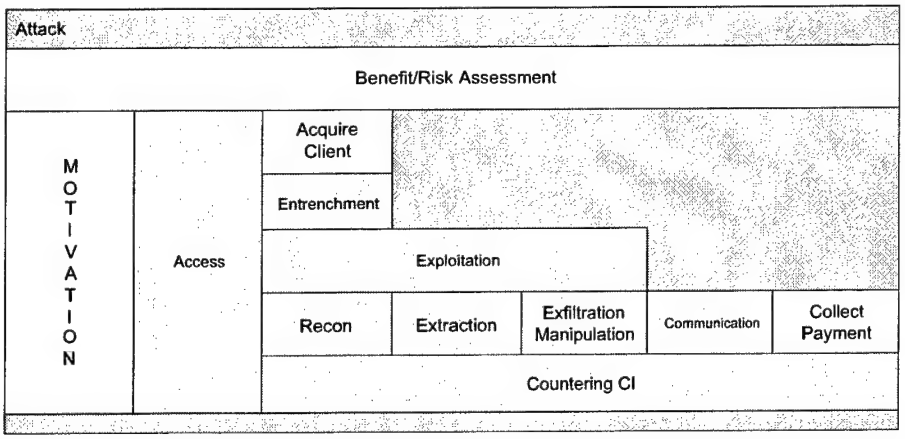

Figure 4.7
Top-Level View of Model

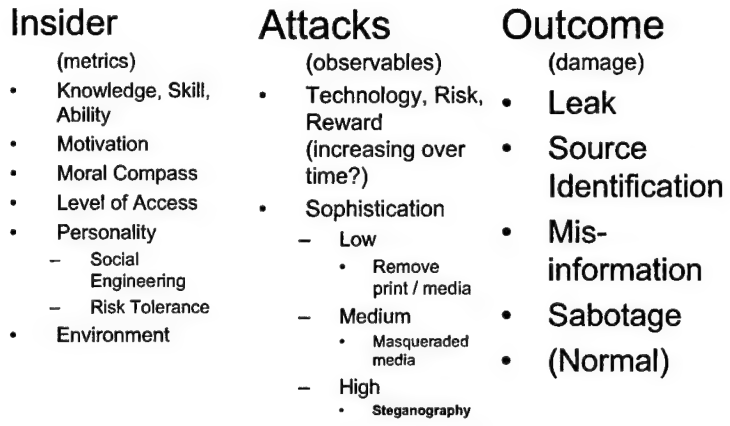

- communication
- manipulation
- counterintelligence.

It could also include other cyber-related activities. The group indicated the level of sophistication for each action: low, medium, or high.

The group then compared this taxonomy with the set of "attack actions" produced by the "Vulnerabilities and Exploits" group (listed in Table 3.1), resulting in the listing shown in Figure 4.8. In this figure, row numbers from Table 3.1 are shown for all exploits from that table. Exploits that do not appear in Table 3.1 are italicized, and seven exploits (numbers 7, 16, 19, 20, 26, 27, 33) in Table 3.1 did not have a place in this group's taxonomy.[3] The parenthesized suffix letters C, I, A, E in Figure 4.8 indicate whether the actions would lead to a breach of information Confidentiality, Integrity, Availability, or would be an Enabler of other attacks.

Definitions

As mentioned earlier, each group was asked to be clear about the definitions of key terms it used. This group listed the following definitions, which cover a number of terms not defined by the IC System Models group (Chapter Two):

- *Threat*: A potential for deliberate attack or an inadvertent compromise of an organization's mission
- *Attack*: A deliberate attempt to compromise an organization's mission
- *Case*: All events and states in the world associated with a related set of attacks. A case may also be a set of other cases.

[3] It would have been desirable to resolve such mismatches among the various taxonomies used by different breakout groups, but the workshop schedule did not permit it. Reconciliation of the various definitions and taxonomies in this report would be a useful follow-on research activity.

Figure 4.8
Insider Actions Taxonomy Cross-Referenced with Vulnerabilities and Exploits List

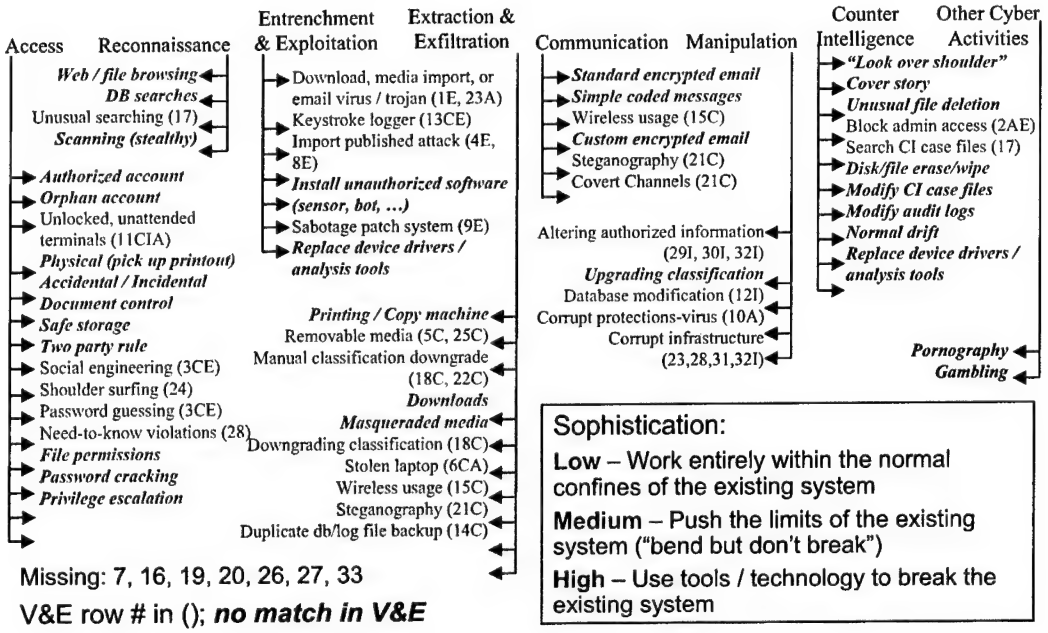

- *Event*: Something that happens in the world. There are atomic events, and logical collections of events can also be events.
- *Model*: An abstract representation of some portion of the world (examples: sensor model, attack model, user model, infrastructure model)
- *Observable*: Any event or state element that can (or could) be measured
- *Sensor*: Measures an observable. Some sensors include Level 1 fusion, and so may also generate a detection.
- *Observation*: The output of a sensor. May be a raw signal, a detection, or both.
- *Detection*: A decision based on processing of observations (or collections of observations)
- *Fusion Level 0*: Raw (signal) output from a sensor
- *Fusion Level 1*: Processed (analyzed) sensor output, often includes an initial detection.

Grand Challenges—Research Issues

The group stated as its overall objective, "We are not building a giant Intrusion Detection System; we are trying to build a set of tools to produce indicators to help the organization investigate anomalous patterns of behavior. We want to shorten the time from defection to detection."

With this in mind, the group listed the following as research issues:

- Extend the taxonomy of insider actions to include, among other things,

  - —mappings from user roles to expected actions and to permitted actions. (The first helps characterize normal behavior; the differences between the two mappings help characterize malicious behavior.)
  - —mappings from actions to observables and observables to sensors. (This should include both existing sensors and desirable new sensors. One approach would be to work with CI experts to identify sensors that would detect the events and extract the information that they use.)
- Develop a characterization of normal system administrator behavior, possibly using Glass Box software (see Appendix H).
- Develop techniques for identifying triggers that distinguish between normal and malicious behavior.
- Develop techniques to identify multiple cases with a common objective.
- Build a library of attacks (scenarios) that researchers can use to train and test their anomaly detection systems.

The first research issue is related to the question of what datasets the researchers need to pursue their research. Some researchers said they could use a realistic document corpus, possibly the NIMD (see Appendix I) WMD corpus. There were several researchers working on methods for anomaly detection, and they all agreed that they need data having sets of attacks interleaved with normal user activity. They need preliminary datasets to train their anomaly detection systems and more sophisticated datasets to validate and tune their systems.

There were five proposals for how to generate the datasets, listed below in approximate order of preference:

A. The existing data from the MITRE workshop might be adequate as a starter set.
B. It might be possible to use existing Glass Box data as a source of normal activity. It would be good to find a way to inject attack data, from the library generated above, into these existing data.
C. An alternative would be to seed the Glass Boxed environment with some attackers in the future. These attackers could mount attacks from the library while other users are going about their normal activities. Glass Box could collect the data and provide them to the researchers.

Both B and C depend on a positive assessment of Glass Box data by the researchers. Right now, it is not known for sure if Glass Box provides what they need, but participants are hopeful.

D. It would be possible to develop a simulator to generate the required datasets. The initial datasets could be fairly simple, requiring a simple simulator. The hope would be that the simulator's capabilities would grow in a manner and at a rate compatible with the expanding needs of the researchers.
E. If none of the above worked out, it would be possible for the researchers to share the datasets they create for themselves.

Whichever approaches are chosen, it will be necessary to define the information needed and a data schema and format so that the researchers can share the datasets.

Surprising Lessons Learned

In concluding its deliberations, the group listed the most surprising outcomes from its discussions as the following:

- Workshop participants still don't have a good understanding of what is important to look for as opposed to what is not (but perhaps they just didn't focus attention on this issue).
- They became aware of the availability of NIMD data and tools, which are seemingly relevant and were previously unknown to them.
- The group was struck by the lack of clear event models.

CHAPTER FIVE

Event Characterization

Terminology

This breakout group[1] started by asking, "What is an *event*?"

- An *event* is an occurrence in a system that is *directly observable.*
- It does not reflect judgment or interpretation (by contrast with detection).
- It is *discrete* (atomic) at the chosen level of abstraction.
- Events are *composable* into higher-level events.
- *Inline* or *mediating sensors* are useful for exposing events in real-time.
- Events can be characterized by data of varying degrees of *fidelity.*

Considering event modeling, there are two kinds of models: models characterizing expected (normal) behavior and malevolent insider (MI) behavior models. A *significant event* fits or deviates from one of these models. That raises the following questions:

- How should we characterize events so that we can determine that they are significant?
- What information needs to be captured when an event occurs so that we can determine if the event fits or deviates from a model?

The set of *relevant events* is influenced by (but should not be limited by) the models being fit.

Events—Considerations

The group then discussed *host-based events* as having the following characteristics:[2]

- Action by a subject involving an object
- Action is then defined as request for or the result of a system service or high-level application activity
  — Can be high or low-level
  — Characterized by time, arguments, return value, status, etc.

---

[1] Participants were Bruce Gabrielson, Greg Kipper, Elizabeth Liddy, Roy Maxion, Kymie Tan, Lisa Yanguas, Dennis Heimbigner, Scott Lewandowski, David Mankins, T. J. Smith, and Feiyi Wang.

[2] The general issue of event characterization has been studied by others. The group used (but was not limited to) past work in this area by Price (1997) and Doyle et al. (2001).

- *Subject* is defined as the user and the process acting on behalf of the user
  - —Characterized by user name, identifiers (e.g., electronic user identification [EUID], radio frequency user identification [RUID], impersonation token, etc.), group membership identifiers, process identifiers, terminal identifiers, etc.
- *Object* is considered to be: a protected system resource
  - —Can be physical or logical
  - —Characterized by names and/or identifiers, type, access permissions, locations, owner, etc. (information available varies greatly)
    - If an action changes the object's attributes, the old and new information should be recorded.

With this terminology in hand, the question then becomes, "How can it be determined if an event is significant?" The group believed that for each event, there are two hypotheses to be tested:

- The event is the result of/indicator of/caused by the activities of a malicious insider (MI), or
- The event is *not* the result of/indicator of/caused by the activities of an MI.

What can be observed to test these hypotheses? The group listed four possibilities:

- Things that we *currently* observe (can and do observe)
- Things we *would like to* observe (can but do not observe)
- Things we *cannot* observe, even in principle (e.g., intent)
- Things we can *indirectly* observe (i.e., inferences from observations).

An event "life cycle" was characterized as

- the recording, or sensing, of events
- interpretation activities—fitting events to models, to determine if they were really caused by an MI and to correlate them (or at least attempt to attribute them to the same insider)
- analysis of "significant" events to learn about the MI: identify his goals, infer the past and predict future activity, and assess damage, both performed and potential.

Data Collection

The group then developed the following "waterfall" diagram (Figure 5.1) to indicate processes that data about an event undergo, and indicating that steps further toward the upper right will lead to increasing probability that the hypothesis being tested about an event will test as "true."

Collection and Analysis

The diagram in Figure 5.1 lists six categories of data. But what processes create one form of data from another? Given an event (the leftmost box in Figure 5.1), the group characterized the collection effort as having the steps shown in Figure 5.2, involving the *actions* of sensors and detectors in converting raw events into indicators or reports (boxes 3 and 4 in Figure 5.1).

Then, given a report, the additional data boxes in Figure 5.1 are created through "fusion" actions or analysis steps (which themselves may access various models as part of their processing), as shown in Figure 5.3.

In this manner (as shown in Figures 5.2 and 5.3), raw events are transformed through a series of processing steps into a case, or set of incidents (the rightmost box in Figure 5.1).

Figure 5.1
Data Collection Steps Regarding an Event

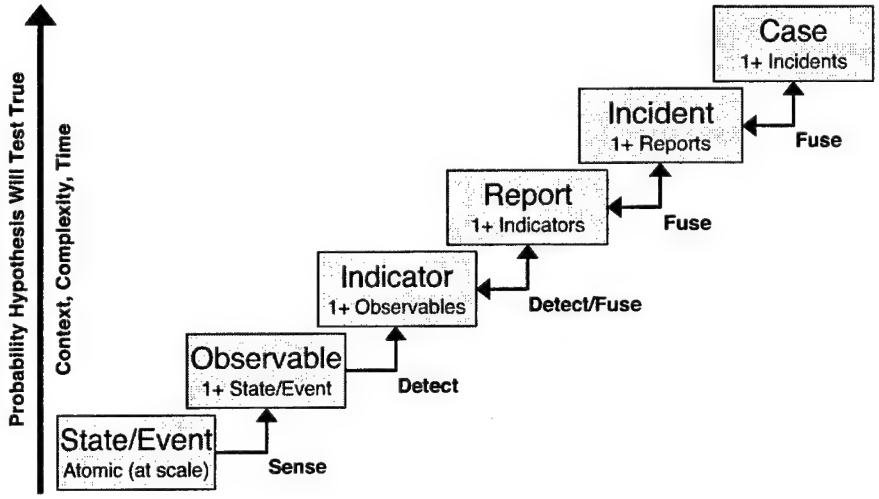

Figure 5.2
Collection Steps

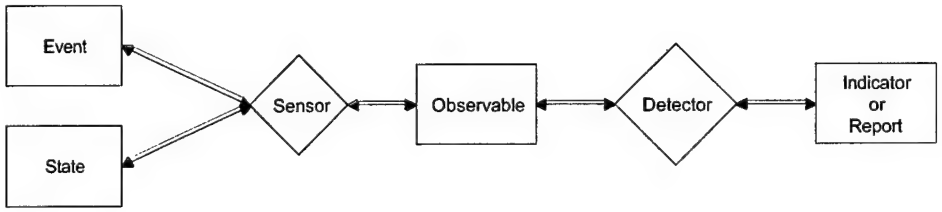

Figure 5.3
Analysis Steps

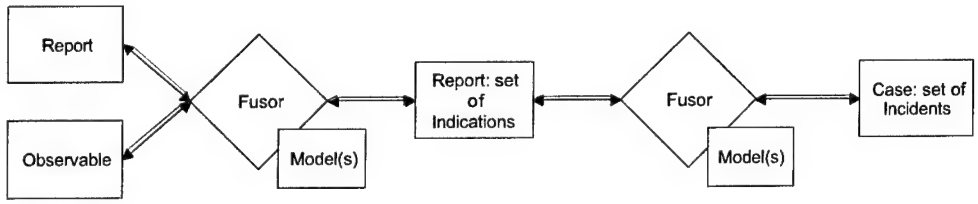

Observables

The group then turned its attention to the observables resulting from events caused by malicious insiders. Its general approach was to create a taxonomy for the insider, including his/her goals, actions to support those goals, and observables associated with the actions.

Then, for each observable, the questions are

- Which are detectable and which are undetectable?
- What new/improved sensors are required to adequately capture the observable?
- What information must be included in reports regarding these observables for the reports to be useful?

The group decided to focus on cyber events only (that is, information that information assurance researchers and operators could provide to counterintelligence personnel), and to disregard most physical observables, even if they can be converted into cyber observables.

In briefing these results, the group felt it had made a good start on this effort, but that more work would be required to complete the task.

Observables from Attacks on Confidentiality

Regarding attacks on the confidentiality of information, the group felt there were two means by which those attacks could be detected:

- Identify data access consistent with the MI threat.
- Detect any exfiltration (with the need for special exceptions for moving data among organizations or among security levels).

Disclosure of confidential information could be performed by a malicious insider through a number of means:

- Accessing legitimately
- Violating "need to know"
- Stealing/using another identity
- Exploiting a misconfiguration
- Exploiting an application vulnerability
- Violating access control (e.g., via privilege escalation).

Malicious insider access to the confidential data could create observables in the following areas:

- Host activity (system calls, resource utilization, etc.)
- Host, operating system (OS), middleware audit/log data
- Host activity records (e.g., browser cache)
- Network activity
- Application records of access (but probably does not apply to exploitation of an application vulnerability).

Exfiltration of confidential data obtained in an exploit could then be achieved by

- writing the data to various media
- media/HW insertion
- network activity (e-mail, scp, ftp, replication, print, etc.)
- context switches
- application-level commands
- host, application, and server print logs
- unobservable techniques: memorize, create images (camera), direct broadcast.

Observables from Corruption of Information

The group indicated that, for automated systems, the distinction between overt/blatant and covert/subtle changes to intelligence data or information is not relevant; it requires human judgment to make that distinction.

The observables that might be detected resulting from adding, changing, or deleting data might be

- data inconsistency (to the extent that consistency can be codified in a model)
- host activity (system calls, resource utilization, etc.)
- host, OS, middleware audit/log data
- host activity records (e.g., browser cache)
- network activity
- application records of access (hopefully characterized in semantic terms, e.g., a global find/replace).

Observables from Degradation of Availability/Access to Information

A third type of attack on intelligence data involves denying others access to critical, relevant information. Access might be degraded or denied to networks, hosts, applications, or by changes in system policy (e.g., as recorded in online distribution and dissemination lists). Observables resulting from attempts at such denial are as follows:

- Degrade networks
  - — Network infrastructure audit/log data
  - — Host, OS, middleware audit/log data (e.g., broken/denied connections) that is *causing* or *resulting from* the denial
  - — Direct network measurement (e.g., tcpdump).

- Degrade hosts
  - —Host, OS, middleware, application audit/log data (e.g., process accounting, general host statistics)
  - —Application indicators (e.g., misbehaving applications).
- Degrade applications
  - —An open challenge
  - —Requires an application specification (may be able to check against a quality of service [QoS] model/guarantee).
- Invoke system policy change or exploit system policy (e.g., cause a user to be locked out)
  - —Host, OS, middleware, application audit/log data
  - —Problem ticket trends.

Observables from Pre-Attack Activities

Finally, the group listed several categories of observables from an attacker's reconnaissance or discovery activities, or activities to obtain access to a system or resource:

- Reconnaissance/discovery
  - —Network activity (e.g., probes)
  - —Host, OS, middleware audit/log data
  - —Honeypots (special role)
- Acquire access or control of a system/resource (lifecycle or post-deployment attacks)
  - —Host, OS, middleware audit/log data.

Research Issues and Questions

This group's "grand challenge" research questions were the following.

Research Issues—Event-Related

- An adequately expressive language for describing and recognizing patterns and events is needed
  - —To facilitate identification/classification of user behavior
  - —To codify tell-tale signs of insider activity
  - —Requires multilevel abstractions and mechanisms
  - —To express uncertainty in characterizing events
  - —Requires features like landmark times, temporal intervals, and temporal and state relationships.
- Examples of event reporting language research efforts exist that might be built upon and instructive, such as
  - —STATL (attack recognition, including sequential, iterative, and conditional events)
  - —CISL (ID reporting; useful lessons learned from CIDF)
  - —CYCL (knowledge representation and reasoning, very general)
  - —P-BEST (Emerald recognition and correlation)
  - —IDMEF (Intrusion Detection Message Exchange Format).

- The implications of varying levels of abstraction must be recognized (and addressed).
  - —There are difficulties to be overcome in unifying/flattening or decomposing event records.
  - —Some detectors (e.g., some data fusion systems) may require a single level of abstraction.
- Can events be associated with MI attack phases or goals?
- What is an effective way to prioritize events, incidents, etc., for human analysis? What information needs to be captured to perform the prioritization and to give the human analysts a good starting point?

Research Issues—Creating Useful Sensors

- What capabilities and attributes should sensors have?
- Some novel sensor ideas are worth exploration:
  - —Honeytokens
    - Simple and relatively easy to deploy
    - Major drawbacks: "bad" (inaccurate) data are introduced into the system
    - Often, may be rejected by certain IC organizations
    - Effectiveness against sophisticated (especially malicious insiders) is questionable
  - —Bloom filters: highly compressed signatures (e.g., for a document)
    - Example: Trace denial of service (DoS) attacks using packet signatures
    - Example: Watch for the signature of a honeytoken in host and network activity.
  - —Social network analysis: Watch for changes in user and system communication patterns.
    - Need models of what analysts should be doing.
  - —Semantic analysis: Watch for changes in a user's "topic of interest."
    - Need to model a user's data flows (e.g., e-mails, queries, etc.) semantically.
- What information needs to be recorded in an event record?
  - —What context is required?
  - —What constitutes a meaningful event record?
- What standards and specifications would be useful for ensuring that sensor reports are accurate and that sensors can be easily integrated into new systems (e.g., a common interpretation of terms, a commonly agreed upon scale for severity and confidence, etc.)?
  - —What standard reporting formats, syntaxes, and methods for recording meaning (semantics) are required?
  - —What does it mean for a sensor to be "self-describing"? Is this a necessary and/or useful capability?
- What can be done to protect sensor systems from attack and ensure sensor outputs are trustworthy?
- How can applications be instrumented to extract useful data?
- What must be done to meet "chain of evidence" requirements?

Research Issues—Sensor Applications

- How can *collaborative processing* (combining evidence by using multiple sensors in concert to gain higher overall confidence in sensor reports) be formalized?

  - —What is required so that differing sensors can participate in a collaborative processing system?
  - —How should individual sensor inputs be weighted?
  - —Can "fusion" be used to create useful synthetic/compound events?
  - —Which techniques are most amenable to scaling to large datasets associated with "low and slow" attacks?
- What sensor data or metadata are required to detect changes in a user's/machine's logical identity (i.e., his/its role)? Can events be matched to role-based access control models "in the background" to detect when users/machines switch roles?
- How can trends be identified and represented in sensor reports?
- How can models and/or policy be used to enhance what has been observed (by refining observation streams via noise reduction)? For instance, if policy disallows port scans, all port scans merit investigation.
- How can false alarms and false positives be identified?
  - —How should detectors be tuned, and what issues affect the tuning?
  - —Tuning must be continuous, and needs to accommodate "model drift" (e.g., changes in traffic, data, and other behaviors) but cannot be exploitable by an adversary.

Research Issues—Building and Working with Models

- How should models be expressed?
  - —How can events be related to those models?
  - —Parameterization of model impacts its performance (e.g., a neural network has learning and momentum constants but users don't understand them).
- How can degree of fit to or deviation from the model be measured?
  - —How much deviation merits treating an event as "significant?"

Research Issues—Testing and Evaluation

- What datasets are required to test sensors?
  - —What metrics are relevant to datasets?
  - —What metadata (e.g., ground-truth) need to be associated with the dataset? Where do these metadata come from?
- What tools are required to manipulate datasets into a form that can be processed?

Research Issues—Miscellaneous

- To support analysis, the state of the subject and object when an event occurred needs to be available—how should this be done?
  - —Explicitly: Record system state data directly in the event record.
  - —Implicitly: Recreate the state of the system from other recorded observations (some of which may be in prior event records).
  - —What are the pros and cons of self-contained versus cumulative audit records? What are the pros and cons of the associated redundancy?
- When should audit records be generated for an action? (Consider non-atomic actions: What are the pros and cons of auditing before and after the action is completed?)

- How should the level of attacker sophistication impact sensor development and research?
- How can sensor information be shared effectively?
  - How can sensors be most effectively used, deployed, and directed? What feedback can be used to help with these tasks?
- How can backward chaining from a hypothetical attack to the events that comprise the early stages of the attack be accomplished in an automated manner?
  - We need to determine if events are indications and warning of that attack at the time they are observed.

Grand Challenge Research Problems

The group concluded its deliberations by creating problem descriptions and research challenges for two "grand challenges."

Challenge 1: Combining Events

Problem Description. Combine events from one or more sensors (possibly of various types and/or different levels of abstraction) to facilitate building systems that

- test hypotheses about MI activity
- gain higher overall confidence in sensor reports
- detect MI activity that is not detectable using a single event record
- reduce data without adversely impacting detection.

Challenge. Determine what is required so that arbitrary sensors can participate in a collaborative processing system. This involves the following:

- Developing a "calculus of evidence" (Does your evidence support your hypothesis?)
- Dealing with metrics (e.g., normalization, new metrics, meta-metrics, finding a common frame of reference, weighting of individual inputs, syntax and semantics of metric reports, dealing with multiple levels of abstraction)
- Identifying and expressing relationships among pieces of evidence
- Determining how "fusion" can be used to create useful synthetic/compound events
- Identifying techniques that can scale to large datasets associated with "low and slow" attacks.

Challenge 2: Exploiting Models and Policies

Problem Description. Improve detection performance by using models and policies.

- Models
  - Run-time: Focus attention on the most significant observations (feature extraction).
  - Design-time: Reason about users and systems to recognize observations that are not currently receiving focus (this can guide development of novel policies and detection methods).

- Policies
  — Run-time: Policy violations are strong indicators of unauthorized behavior.
  — Design-time: Direct the deployment of sensors (to observe all policy violations).

Challenges.

- How should models and policies be expressed, monitored, measured, etc.?
- How can events be related to models and policies?
- How can machines (software programs) understand the implications of policies?
- How do we design models that are usable by "regular users?"
- What sensor data and metadata are required to detect changes in a user's/machine's logical identity or role?

APPENDIX A

Workshop Invitation

February 2, 2004

Dear Colleague,

Public concern for information security has been focused on the problem of preventing harm resulting from the actions of an outside attacker. However, there is greater risk of damage due to hostile actions performed by insiders, possibly in collaboration with agents of an outside organization. One insider may transmit large volumes of sensitive information outside the organization with little risk of detection. Another insider might make subtle, or not so subtle, modifications to a critical database, resulting in large losses in human or financial terms. These behaviors are especially insidious if the insider installs software that could perpetuate this behavior even after the attacker has left the organization.

The risk of insider attacks is greatest for systems that contain high value, mission critical data. These high value targets may be classified or unclassified Government systems or systems in the private sector. They all tend to attract the attention of well-funded organizations, including foreign governments, that are willing to recruit insiders in an effort to mount sustained, well-planned penetrations into an organization's cyber assets.

The sorts of attacks mounted by these adversaries depend on many factors, including motivation and objectives, level of knowledge about the system, authorized access to the system, tolerance for risk, and specific computer skills, including the skills of outside collaborators working with the insider. Effective defense against insider attacks must be based on a realistic understanding of the behaviors that different classes of insiders are likely to exhibit.

ARDA invites you to participate in a workshop on "Understanding the Insider Threat." This three day workshop will bring together members of the Intelligence Community (IC) with specific knowledge of IC document management systems and IC business practices, individuals with knowledge of inside attackers, both within and outside the IC, and researchers who are involved in developing technology to counter insider threats. The objectives of the workshop are:

- To generate and capture domain knowledge that will benefit the broad base of researchers studying the Insider Threat. This includes but is not limited to knowledge about:
  - —Inside attacker characteristics, including the vulnerabilities they tend to exploit, and the attack methods they use,
  - —Attack characterization, including the necessary or likely pre-conditions for an attack, the observables generated during an attack, and the effects of the attack.
  - —The network and application systems used by the IC for document management, including the mechanisms used to protect the systems and data.
  - —IC business models for generating and controlling access to documents.
- To foster cooperation among researchers by developing, to the extent it is practical, methods for describing common aspects of their work, such as event characterization, attack and attacker classification, etc.
- To focus researchers on specific systems and problems of interest to the IC. We expect these to take the form of challenge problems.

We are looking forward to a fruitful discussion on these important topics. ARDA intends to document the results of the workshop in a RAND report.

Enclosed is additional information regarding the workshop format, logistics, and some important pre-work we would like each of you to do prior to the meeting.

Workshop Location and Dates:

The workshop will be held 2–4 March 2004 at McAfee Research, Network Associates, Inc., 15204 Omega Drive, Suite 300, Rockville, MD 20850. Office phone: 301-527-9500. See below for directions and hotel information.

Workshop Format:

As you can see from the appended agenda, this workshop will be less presentation-centric than many workshops. The approach is to have presentations by domain experts to provide background material and to then break into subgroups to discuss issues and approaches, followed by plenary sessions to compare results and synchronize the groups. When registering, please provide a prioritized list of the subgroup topics (identified in the agenda below) ranked according to your interests and expertise. We will use these lists to pre-assign attendees to breakout sessions.

Pre-Workshop Material:

Prior to the workshop we will distribute a package of reading material to help you prepare. We request that each of you please prepare (no more than 2 pages) some informal thoughts on the following topics as related to your area of expertise:

i. A description and assessment of the top two or three issues associated with understanding the Insider Threat,
ii. Your thoughts on how to address these issues, and
iii. Very brief descriptions of related efforts being pursued in your area of expertise.

Registration:

By Feb 10, please e-mail your registration information, or regrets, to **research@mcafeesecurity.com.** Include the following details:

- Full name
- Prefix (Mr. / Ms. / Dr. / military rank / etc.)
- Name for badge
- Organization
- Postal address
- E-mail address
- Telephone number
- Fax number
- Cell number
- Hotel (if applicable)
- Your prioritized list of subgroup topics ranked according to your interests and expertise

Meals / Refreshments:

Continental breakfast, lunch, afternoon snack, and drinks will be provided each day.

Workshop Fee:

There is a workshop fee of $75.00 to cover cost of meals, refreshments, and workshop supplies. By February 17, 2004, print a copy of your registration information and mail it, along with check or money order payable to **Network Associates** for your workshop fee, to Dana Coon, McAfee Research, 15204 Omega Drive, Suite 300, Rockville, MD 20850.

Information on Hotels:

There are rooms generally available at a wide variety and number of hotels in the area. Please make your reservations as soon as possible. Here are a few suggestions:

Homestead Gaithersburg/Rockville
2621 Research Blvd.
Rockville, MD
Tel: 301-987-9100

Courtyard by Marriott–Rockville
2500 Research Blvd.
Rockville, MD
Tel: 301-670-6700

SpringHill Suites Gaithersburg
9715 Washingtonian Blvd.
Gaithersburg, MD 20878
Tel: 301-987-0900
Van service available

Gaithersburg Marriott Washingtonian Center
9751 Washingtonian Blvd.
Gaithersburg, MD 20878
Tel: 301-590-0044
Rooms have not been blocked.

Directions to McAfee Research:

McAfee Research is located in the Rockville, MD, offices of Network Associates, in the greater Washington, DC metropolitan area, off Interstate 270.

From I-270 NORTH, take Exit 8 for Shady Grove Road WEST. Stay in the right lanes on the frontage road until it veers right, then stay in one of the left 2 lanes and proceed to the light. Make a LEFT onto Shady Grove Road. Cross over I-270. Turn RIGHT onto Research Boulevard. Turn LEFT onto Omega Drive. The Network Associates building is on the right.

From I-270 SOUTH, take Exit 8 for Shady Grove Road WEST. Stay in the right lanes on the frontage road until it veers right, and then stay in the right lane. Take the exit for Omega Drive in the far right lane. (If you miss the Omega Drive exit, you can still pick up the directions from I-270 NORTH above, starting with the point where you cross over I-270.) Turn LEFT onto Omega Drive. The Network Associates building is on the right.

Turn RIGHT into the entrance. The Network Associates building is now to your right front. A 4-foot cube-shaped sign with the Network Associates logo and name is in front of the building.

Enter the building into the lobby area through the double glass doors. Take an elevator located on your left to the third floor.

Parking

Parking is in the areas in front of, to the right of, and behind our building. You may ignore signs indicating permit parking. The only restrictions are marked handicapped spaces.

Shuttle / Van Service

If you wish to use this service, contact your hotel desk for availability and procedure. Note that there may be a fee.

Contact Information—For more information, please contact:

Technical:
Dick Brackney: 301-688–7092
John Farrell: 443-479-4370

Administrative/Logistics:
Jack Oden: 301-947-7159, Cell: 703-402-8574
Dana Coon: 301-947-7275.

APPENDIX B

Workshop Agenda

ARDA Understanding the Insider Threat Workshop Agenda

The general approach for the workshop will be to use the large group to provide background information to the participants, focus the group, and gather results, and to use a set of small-group breakout sessions to discuss the issues and recommend solutions or areas of needed research.

2 March 2004

The objective for this day is to establish a baseline context and knowledge set for you to use for the rest of the workshop. Speakers with domain knowledge will provide introductory presentations, and each of you will have the opportunity to state your priorities and to describe what you have to offer the group.

7:30 AM Continental Breakfast/Coffee

Session 1: Workshop Purpose and Context—Plenary Session

8:00 Welcome and Workshop Charge—Dick Brackney
8:15 Introductory Remarks—Sherrill Nicely, Information Assurance Director, IC CIO
8:30 Speaking with Analysts: Observations of Current Practices with Massive Data, William Wright, Oculus Info Inc.
9:10 IC Document Management and Dissemination Systems, speaker TBD
10:00 Break
10:30 Overview of the Hanssen case, Robert Anderson, the RAND Corporation
10:50 Insider Behavior—Speaker TBD
11:30 Report on the MITRE Workshop on Indications and Warnings for the Insider Threat: Mark Maybury, The MITRE Corp.

12:15 Lunch

Session 2: Identify Needs and Research Topics—Breakouts

1:30 Review agenda, expectations, and first breakout assignment
Break into subgroups: IC System Models, Attacker Models, Event Characterization, Vulnerabilities and Exploits
1:45 Review the pre-work (use to scope the issues)
Each Group Participant Lists (< 5 minutes each):
— Knowledge acquired to date
— Top needs
— Most promising sources and approaches for acquisition
— Summarize key points in Power Point file
3:30 Break
4:00 Report out to large group (15 minutes per group)
— Each group presents key points to large group
— Discuss for clarity
— Consolidate points where possible

3 March 2004

The objective for today is to identify, generate, and capture the knowledge that researchers need to work in this area. We expect this to happen through a dialogue between researchers, domain experts, and security practitioners. Tomorrow you will use this knowledge to define a set of challenge problems for the researchers.

7:30 AM Continental Breakfast/Coffee

Session 3: Generating the Knowledge

8:00 Review Agenda, expectations, and second breakout
Break into subgroups: IC System Models, Attacker Models, Event Characterization, Vulnerabilities and Exploits
8:15 Identify and address top knowledge needs
— Identify what is known: document sources
— Identify what else researchers need to know
- Work as team to fill knowledge gaps
- Identify approaches for further work

— Summarize key points in Power Point file
10:15 Break
10:45 Large group report out (15 min per group)
— Share results with large group & discuss
— Integrate results across groups

12:00 Lunch

Session 4: Organizing the Knowledge

1:15 Reconvene in Large Group: Review Agenda, expectations, and third breakout

Break into subgroups: IC System Models, Attacker Models, Event Characterization, Vulnerabilities and Exploits

1:30 Organize the knowledge identified in the AM session in a manner usable by researchers
- —Agree on principles for organization
- —Organize the knowledge
- —Identify and prioritize gaps
- —Fill in gaps where feasible

3:30 Break

4:00 Large group report out (15 min per group)
- —Share results with large group & discuss
- —Integrate results across groups

5:00 Adjourn for day

4 March 2004

The objective for this day is to define a set of research problems that are of interest to practitioners and to identify useful intermediate results that we can use to measure progress.

7:30 Continental Breakfast/Coffee

Session 5: Defining Measures of Success

8:00 Review results so far and set agenda for Day 3
Break into subgroups

8:30 Identify research problems whose solution will benefit the security practitioners
- —Define the problems
  - Brainstorm, then pick the top two
  - Use the knowledge base developed yesterday
- —Explain the benefits to the practitioners of solving the problems
- —Identify useful partial results and their benefit to the practitioners
- —Identify remaining knowledge gaps and suggest ways to fill them
- —Summarize key points in Power Point file

10:30 Break

11:00 Large group report out (15 min per group)
- —Share results with large group & discuss
- —Integrate results across groups

12:00 Lunch

1:00 Wrap-up
- —Summary of results
- —Final opportunity to comment
- —Feedback from ARDA
- —Burn CD's and distribute to the participants

2:30 Adjourn

APPENDIX C

Links to Read-Ahead Materials

All workshop participants were provided links to the following materials, which were deemed relevant as background and "read-ahead" material.

Reports

- DOD Insider Threat Mitigation: Final Report of the Insider Threat Process Team (http://www.c3i.osd.mil/org/sio/iptreport4_26dbl.doc)
- RAND Insider Threat Report (http://www.rand.org/publications/CF/CF163/)
- SecurityFocus Newsletter 132 (http://www.securityfocus.com/infocus/1546/)
- References in Characterizing Threat Paper
- White Paper: Cyber-Security and the Insider Threat to Classified Information, November 1–2, 2000, Computer Science and Telecommunications Board (CSTB) (http://www7.nationalacademies.org/CSTB/whitepaper_insiderthreat.html)
- Inside the Mind of the Insider (http://www.securitymanagement.com/library/000762.html)
- Espionage Against the United States by American Citizens 1947–2001 (Defense Personnel Security Research Center) (http://www.ncix.gov/news/2002/oct/Espionage.pdf)
- The Insider Threat to U.S. Government Information Systems (NSTISSAM INFOSEC) (http://www.nstissc.gov/Assets/pdf/NSTISSAM_INFOSEC1-99.pdf)
- R. Anderson, R. Brackney, T. Bozek, Advanced Network Defense Research: Proceedings of a Workshop (CF-159-NSA) (http://www.rand.org/publications/CF/CF159/CF159.pdf).

Collections of Cases and Reports

- Recent Espionage Cases 1975–1999 (Defense Security Service) (http://www.dss.mil/training/espionage/)
- DSS Employee Security Training
  — Spy Stories (http://www.dss.mil/training/csg/security/Spystory/Intro.htm)
  — Treason 101(http://www.dss.mil/training/csg/security/Treason/Intro.htm)
- Federation of American Scientists Counter-Intelligence Operations (http://www.fas.org/irp/ops/ci/)

- CI Centre Counterintelligence Reference Materials (http://www.cicentre.com/LINKS_Reference_Material.htm)
- Office of the National Counterintelligence Executive (NCIX) (http://www.ncix.gov/news/index.html).

Cases

- Robert Hanssen (FBI)
  - —A Review of the FBI's Performance in Deterring, Detecting, and Investigating the Espionage Activities of Robert Hanssen (Aug 14, 2003) (http://www.usdoj.gov/oig/special/03-08/index.htm)
  - — CI Centre Article (http://www.cicentre.com/Documents/DOC_Hanssen_1.htm)

  Webster Commission Report (http://www.fas.org/irp/agency/doj/fbi/websterreport.html)
  - —Attorney General Webster Commission Report Commentary (same as above)
  - — Hanssen Computerworld article (http://www.computerworld.com/securitytopics/security/story/0,10801,57889,00.html)
  - —Webster Commission CI Center article (http://www.cicentre.com/Documents/DOC_Quotes_Webster_Report.htm)
- Moonlight Maze
- Ana Montes (DIA)
  - — CI Centre Article (http://www.cicentre.com/Documents/DOC_Montes_1.htm)
- Harold James Nicholson (CIA)
  - —Affidavit (http://www.fas.org/irp/offdocs/nicholson.htm)
- Brian Reagan (NRO)
  - — CI Centre Article (http://www.cicentre.com/Documents/DOC_Regan_1.htm)
  - —Washington Post article on discovery of documents.

Conferences

- Mitigating the Insider Threat: Proceedings of RAND August 2000 Workshop (http://www.rand.org/publications/CF/CF163/)
- IATF Forum (http://www.iatf.net/)
- Practical CounterIntelligence Conference
- PACOM IA Conference (http://www.iaevents.com/)
- 1999 SRI Report (http://www.csl.sri.com/users/neumann/insider-misuse/ins.pdf).

APPENDIX D

Workshop Participants

Robert H. Anderson	RAND Corporation
Richard Brackney	ARDA
Philip Burns	Computer Technology Associates
Matthew Downey	Syracuse Research Corp
Jeremy Epstein	webMethods
Paul Esposito	Defensive Computing Research Office
John Farrell	ARDA
Dana Foat	Defense-wide Information Assurance Program
Bruce Gabrielson	Booz Allen Hamilton
Chris Geib	Honeywell Labs
Joseph Giampapa	Carnegie Mellon University
Alexander Gibson	Battelle Northwest Labs
Terrance Goan	Stottler Henke Associates
Frank Greitzer	Battelle Northwest Labs
Tom Haigh	Adventium Labs
Steven Harp	Adventium Labs
Dennis Heimbigner	University of Colorado
Thomas Hetherington	Applied Research Lab, University of Texas, Austin
William Hunteman	U.S. Department of Energy
Peter Jobusch	Intelink Management Office
Clarence Jones	NSA
Steve Karty	NCS
Kevin Killourhy	Carnegie Mellon University
Greg Kipper	MITRE Corporation
Linda Kiyosaki	NSA
Stephen Laird	Lockheed Martin Orincon
Vincent Lee	USG/CIO
Van Lepthien	University of Colorado
Scott Lewandowski	MIT Lincoln Laboratory

Elizabeth Liddy	Syracuse University
Tom Longstaff	CERT Coordination Center
David Mankins	BBN Technologies
Sara Matzner	Applied Research Labs, University of Texas, Austin
Roy Maxion	Carnegie Mellon University
Mark Maybury	MITRE Corporation
Mark Morrison	MITRE Corporation
Richard Neely	Computer Technology Associates
Sherrill Nicely	CIA
Lucille Nowell	ARDA
Michael Pelican	Honeywell Labs
Marisa Reddy	U.S. Department of the Treasury
David Sames	McAfee Research
Thomas Shackelford	George Mason University
T.J. Smith	MCNC Research & Development Institute
Frederick Steinheiser	U.S. Government
Greg Stephens	MITRE Corporation
Kymie Mei Chen Tan	Carnegie Mellon University
Roshan Thomas	McAfee Research
Shambhu Upadhyaya	University of Buffalo
Feiyi Wang	MCNC Research & Development Institute
Brad Wood	BBN Technologies
Bill Wright	Oculus Information
Edward Wright	Information Extraction & Transport
Lisa Yanguas	NSA/R6

APPENDIX E

Presentation: The Robert Hanssen Case: An Example of the Insider Threat to Sensitive U.S. Information Systems

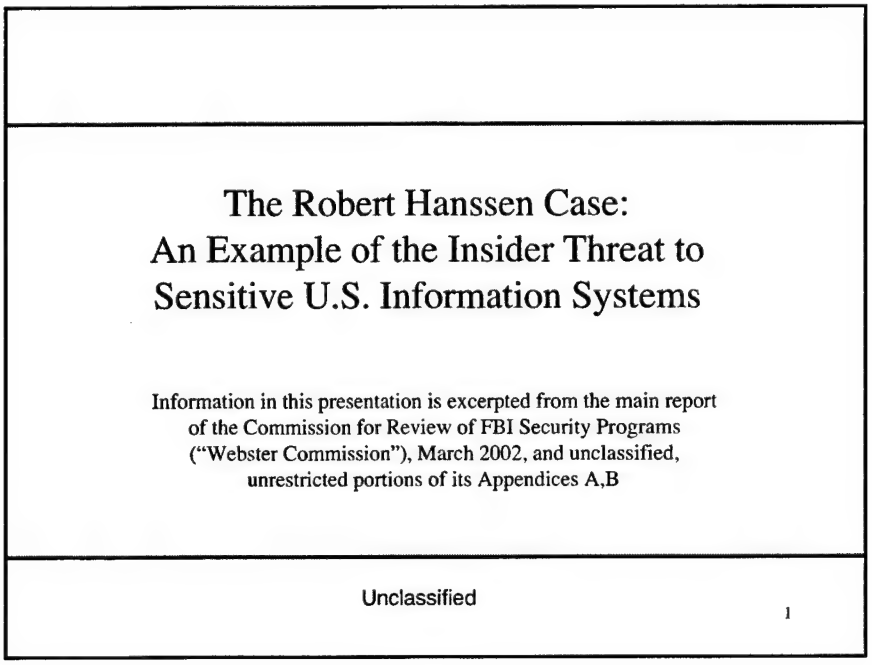

Outline of presentation

- Background
- What did he do?
- How did he do it? What might be done?
- How could he? - problems in FBI infosec
- Relevant FBI systems and info architectures
- Commission recommendations on solution strategies
- Some possible conclusions for this workshop

2

Background

- Robert Hanssen, at the time of his arrest, was an FBI Supervisory Special Agent
- His treason is called "possibly the worst intelligence disaster in US history"
- Over 22 years, he gave the Soviet Union and Russia vast quantities of documents and computer diskettes filled with national security information
- This insider treason is part of a recurring pattern:
  - "Since the 1930s, every US agency involved with national security has been penetrated by foreign agents, with the exception of the US Coast Guard"
  - 117 American citizens have been prosecuted for espionage between 1945-1990 (or there is clear evidence of their guilt). Money appears to be the main factor; most spies volunteered their services. Prominent examples:
    - *Aldrich Ames*, CIA counterintelligence officer (9 years as spy)
    - *Ronald Pelton*, former intelligence analyst for NSA
    - *Jonathan Pollard*, military intelligence analyst, gave Israel 800 classified documents, 1000 cables
    - *John Walker*, retired naval officer, with son and brother, supplied the Soviets with cryptographic material

3

What did he do?

- Downloaded large quantities of information from the main FBI Automated Case Support system
- Searched the Bureau's systems to see whether the FBI had identified his locations as drop sites
- Searched for his name in the system to see if he was the subject of an investigation
- Installed unauthorized software on his office computer
- Hacked into the computer of a Bureau colleague
  - "...purportedly to demonstrate security weaknesses in the computer system." Although discovered, he was not punished for this.
- Photocopied documents at the Bureau, and walked out with them

*
- Walked into classified meetings uninvited ("habitually")
- Visited former colleagues, discussed classified information
- Borrowed a TS/SCI document, photocopied it in the back seat of his car, then returned it

* "Non-technical methods"

4

How did he do it? What might be done?

- With one exception*, all his activities involved:
  - technical access for which he was authorized, or
  - "non-technical methods"
- What might be done (today)?
  - USB "thumb drives" -- a gigabyte on a keychain
    - Will all USB ports be disabled on desktop and laptop computers? Could they be?
  - CD, DVD -- could all CD, DVD write access be disabled?
  - Wireless transmission
    - A wireless card slipped into a laptop or desktop PC, transmitting to an external "base" computer (e.g., in a parked car)
    - Bluetooth local transmission to a cellphone, PDA, etc.
    - Infrared transmission to a local device (e.g., PDA)
    - Digital camera built into a cellphone or PDA
    - Audio recording within a PDA or cellphone

* "Hacking" into a supervisor's computer (term undefined in report)

5

How could he? Systemic problems in FBI InfoSec

- Bureau failed to develop an effective strategy to identify and protect critical information
- Classified information was moved into systems not properly accredited for it
- Until recently, FBI didn't certify and accredit most of its computer systems, including those handling classified information
- Inadequate physical protections
- Lack of adequate documented INFOSEC policies
- FBI failed to ascertain security requirements of "owners" of information, and identify threats and vulnerabilities that must be countered
- Users lack sufficient guidance about critical security features
- FBI failed to limit user access to systems and databases
- Many key InfoSec positions remained unfilled; when filled, staff have inadequate time, support, and authority
- Some FBI systems have insufficient resources to perform required audits; when audits are performed, audit logs are reviewed sporadically, if at all

6

FBI systems architecture (excerpts)

- The FBI operates between 30-60 classified systems (A7). In many cases, the boundaries and missions of these systems were difficult to ascertain (A6)
- Automated Case System (ACS) is the main investigative system of records
  - This was exploited "almost exclusively" by Hanssen during his last period of spying
  - Has access restrictions by office (O) or by list of persons authorized (P)
    - But HQ personnel could access all O cases opened by any field office (A29)
  - Highly sensitive information can be found in unrestricted admin files, a fact that Hanssen exploited (A30-31)
  - After 9/11, HQ mandated that no ACS case may be restricted or deliberately not uploaded without approval of an Asst. Director; HQ later removed additional ACS case restrictions (A35)
  - No formal procedure for terminating accounts when a user's need for access ended

Note: Symbols such as (A7) provide appendix and page reference

7

FBI systems architecture (excerpts, cont.)

- Trilogy is a 3-part system and network upgrade underway
  - But it is being rushed to completion, with security features playing "catch up"
- FBINet
  - A Secret-high FBI network, but TS/SCI data has been processed on it (A42)
- Diskettes
  - Many unmarked, others marked unclassified -- were used on Internet terminals without screening information on the diskettes (A44)
- HQ Data Center Facility
  - Contains a trove of classified information; has backup tapes
  - In a SCIF, staffed 24 hrs/day. All staff polygraphed. All staff cleared for all SCI compartments
  - "If one of these persons were to smuggle these tapes out of the FBI ... the entire content of the FBI's Investigative Mainframe, including ACS and the ASSET database, could be restored outside the FBI and the tapes returned before anyone noticed them missing." (A48)

8

FBI systems architecture (excerpts, cont.)

- Audit logs and trails
  - Were used extensively to trace Hanssen's activity, *after the fact*
  - Investigators able to determine which ACS files Hanssen viewed, how he searched for them, how long he viewed them, and whether he downloaded them (A58)
    - But not whether he did screen dumps while viewing them
  - FBI audit trails were never reviewed in real time
  - Other than occasional and ill-defined reviews, no one at FBI reviewed and analyzed audit data proactively

9

Commission recommendations re. solution strategies

- Countermeasure objectives should be: (B5)
  - Reduce the time between defection and detection
  - Reduce the number of defectors
  - Reduce the number of information compromises (attacks) by a defector, and
  - Reduce the amount of damage caused by each compromise (or attack)
- Distinguish three types of "moles:" (B8-10)
  - *Standard* (has regular user privileges)
  - *Privileged* (more access and privileges (either mission-privileged, or over-privileged)
  - *Cracker* (bends or violates assigned permissions and privileges)

10

Commission recommendations re. solution strategies (cont.)

- Consider requirements of 18 baseline INFOSEC categories (from NSA's INFOSEC Assessment Methodology (B20-96)

INFOSEC documentation	Contingency planning
INFOSEC roles and responsibilities	Maintenance
Identification and authorization	Configuration management
Account management	Labeling
Session controls	Media controls
Networking/connectivity	Physical environment
Telecommunications	Personnel security
Auditing	Education, training, awareness
Malicious code protection	System assurance

11

Some possible conclusions for this workshop

- It may be shortsighted to consider just one "test" system, such as Intelink
  - when the real world comprises a congeries of systems linked in various ways, with different authorities, owners, procedures, criteria -- even in one agency
- Many violations used "nontechnical means"
  - Can technical means aid in discovering these violations?
- If users are unaware/untrained/unappreciative of security markings, restrictions, controls -- data could be available to "normal" violators, not needing extraordinary means
- It is ever more trivial to record large amounts of data on tiny devices, for physical removal -- often with no trace or audit log record
- Audit logs are only as good as the uses to which they're put
  - Too many "false positives", too many resources required to operate and review, too many differing systems' logs not combined, too much data -- all weaken their effectiveness

12

APPENDIX F

Presentation: Overview of the Results of a Recent ARDA Workshop on Cyber Indications and Warning[1]

DSS

Cyber Indications & Warnings Insider Threat Workshop Overview

Dr. Mark Maybury
MITRE

2 March 2004

MITRE

ITC

[1] This presentation is copyrighted by MITRE Corporation. Reproduced by permission.

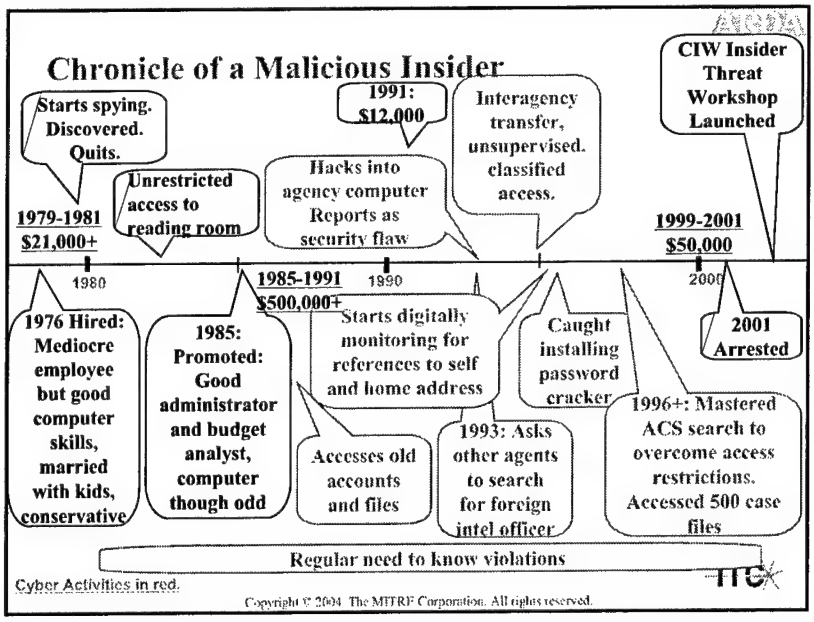
Chronicle of a Malicious Insider
Starts spying. Discovered. Quits.
1979-1981
$21,000+
Unrestricted access to reading room
Hacks into agency computer Reports as security flaw
1991: $12,000
Interagency transfer, unsupervised. classified access.
CIW Insider Threat Workshop Launched
1999-2001
$50,000
1980
1985-1991
$500,000+
1990
2000
1976 Hired: Mediocre employee but good computer skills, married with kids, conservative
1985: Promoted: Good administrator and budget analyst, computer though odd
Starts digitally monitoring for references to self and home address
Caught installing password cracker
2001 Arrested
Accesses old accounts and files
1993: Asks other agents to search for foreign intel officer
1996+: Mastered ACS search to overcome access restrictions. Accessed 500 case files
Regular need to know violations
Cyber Activities in red.
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

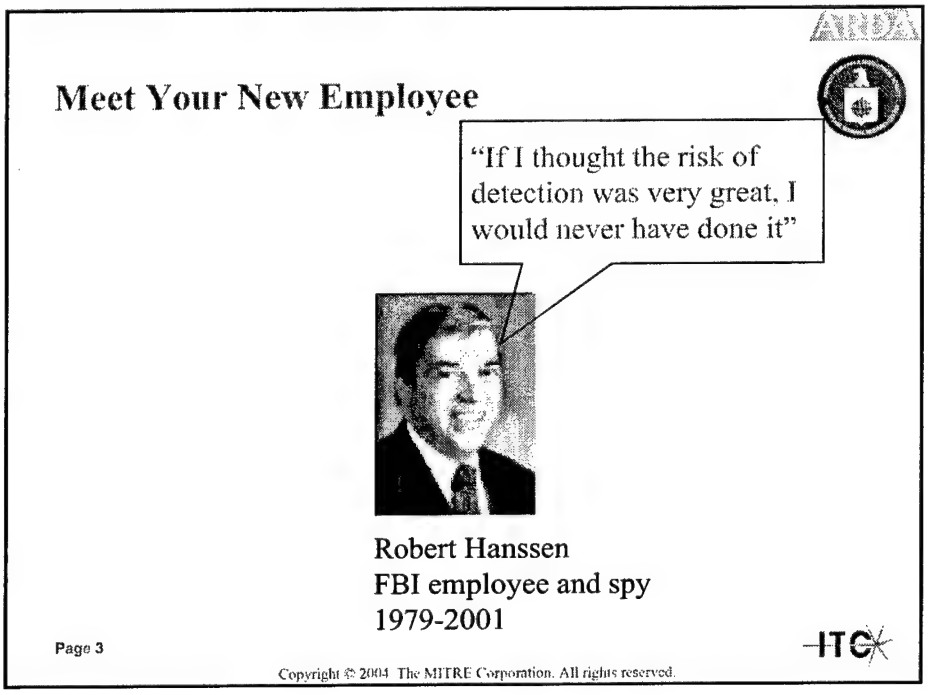
Meet Your New Employee
"If I thought the risk of detection was very great, I would never have done it"
Robert Hanssen
FBI employee and spy
1979-2001
Page 3
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

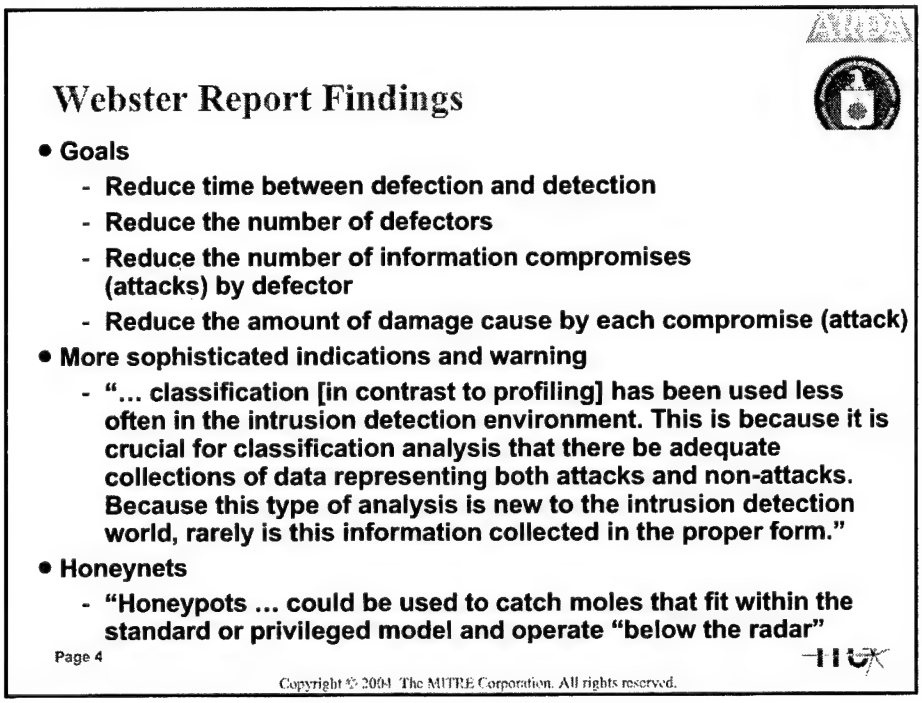
Webster Report Findings
• Goals
- Reduce time between defection and detection
- Reduce the number of defectors
- Reduce the number of information compromises (attacks) by defector
- Reduce the amount of damage cause by each compromise (attack)
• More sophisticated indications and warning
- "... classification [in contrast to profiling] has been used less often in the intrusion detection environment. This is because it is crucial for classification analysis that there be adequate collections of data representing both attacks and non-attacks. Because this type of analysis is new to the intrusion detection world, rarely is this information collected in the proper form."
• Honeynets
- "Honeypots ... could be used to catch moles that fit within the standard or privileged model and operate "below the radar"
Page 4
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

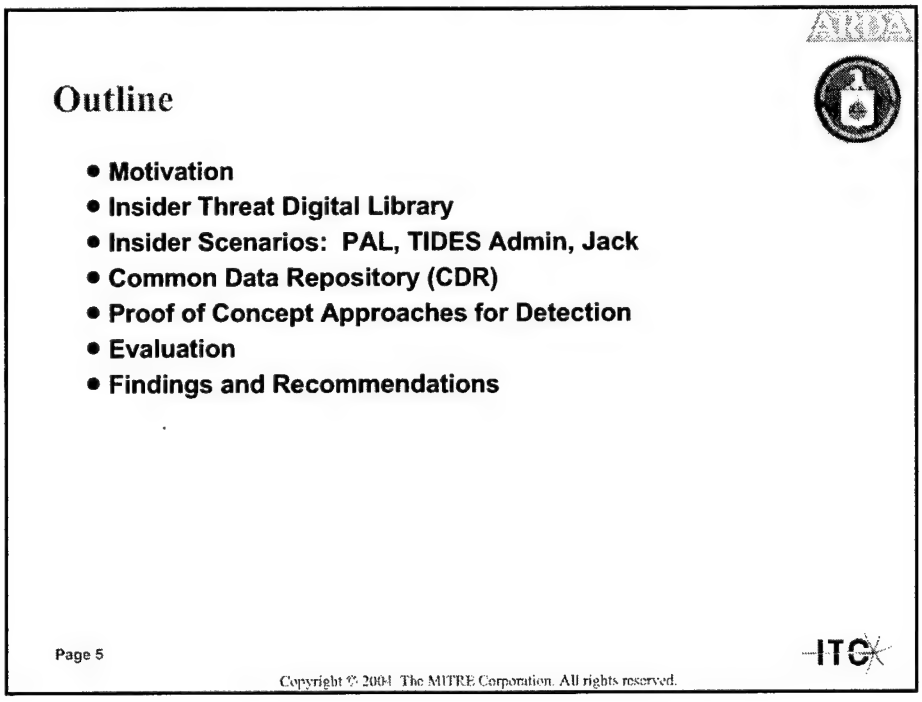
Outline
• Motivation
• Insider Threat Digital Library
• Insider Scenarios: PAL, TIDES Admin, Jack
• Common Data Repository (CDR)
• Proof of Concept Approaches for Detection
• Evaluation
• Findings and Recommendations
Page 5
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

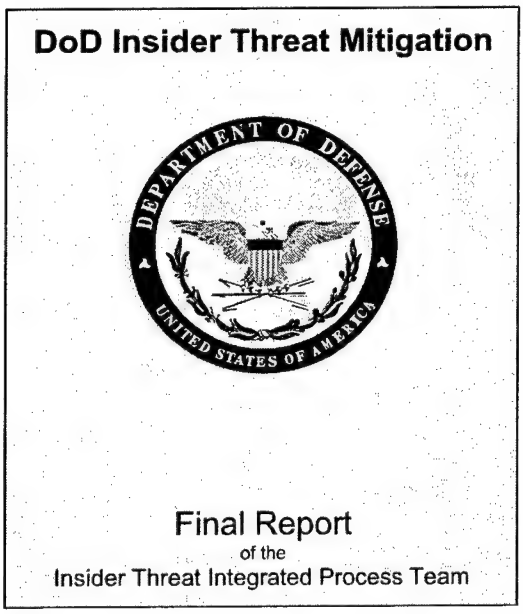

"A recent DoDIG report indicates that, for [over 1,000] investigations, 87 percent of identified intruders into DoD information systems were either employees or others internal to the organization"

Cyber Indicators Exist: Nearly half of suspicious contact reports made to the Defense Security Service by defense contractors begin with an email request for information, especially by foreign organizations, per Gene Smith, a DSS counterintelligence analyst

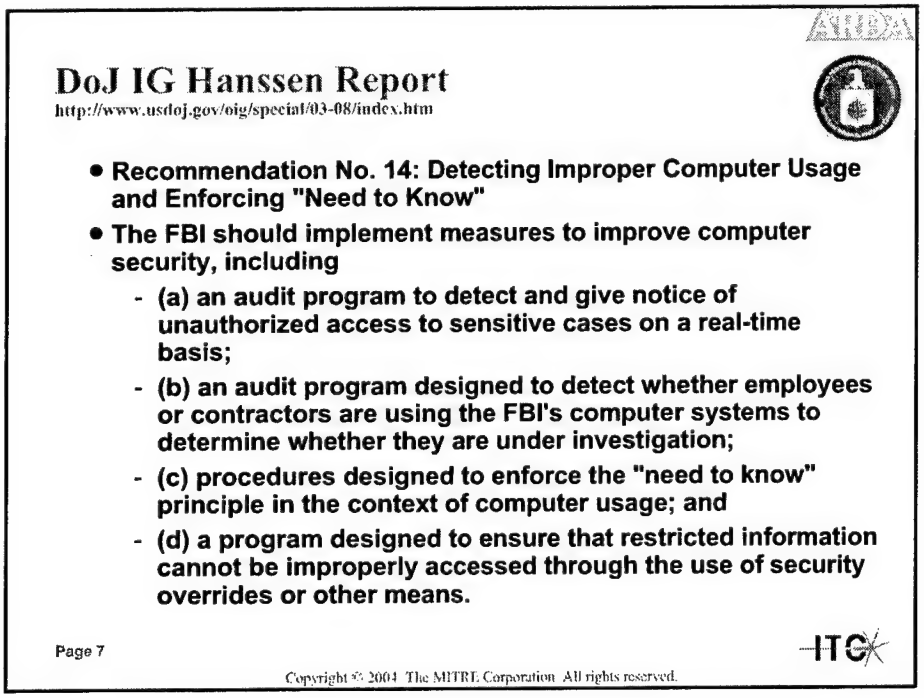

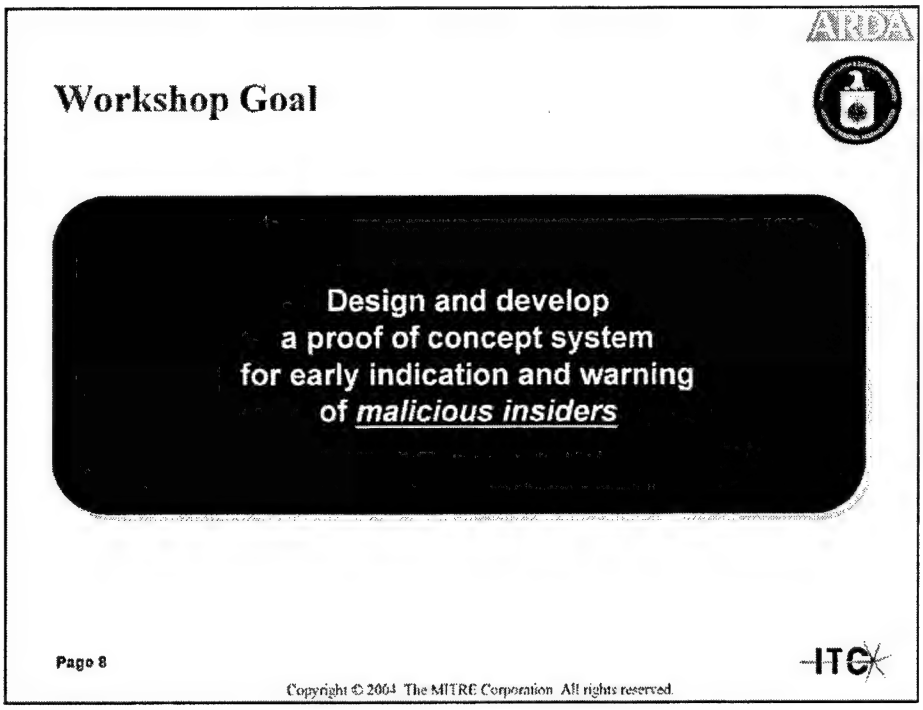

Definitions

- Insider – **Anyone with access, privilege, or knowledge of information systems and services**
- Malicious insider (MI) – **motivated to intentionally adversely impact an organization's mission (e.g., deny, damage, degrade, destroy).**
- Observable - **Anything that can be detected with current tech.**
- Sensor – **Measures an observable (e.g., login, print, delete)**
  - **Sensor logs – recording of observables**
  - **Sensor stream – series of observables from one sensor**
- Indicator – **Identifiable event based on sensor output logs**
- Detect – **Determines event based on processing indicators**
- Report – **Indications and warnings of malicious insider behavior**
- Incident – **Related set of events**
- Fusion – **Processing multiple sensor outputs to provide an enhanced result (e.g., more abstract/concrete, higher confidence)**

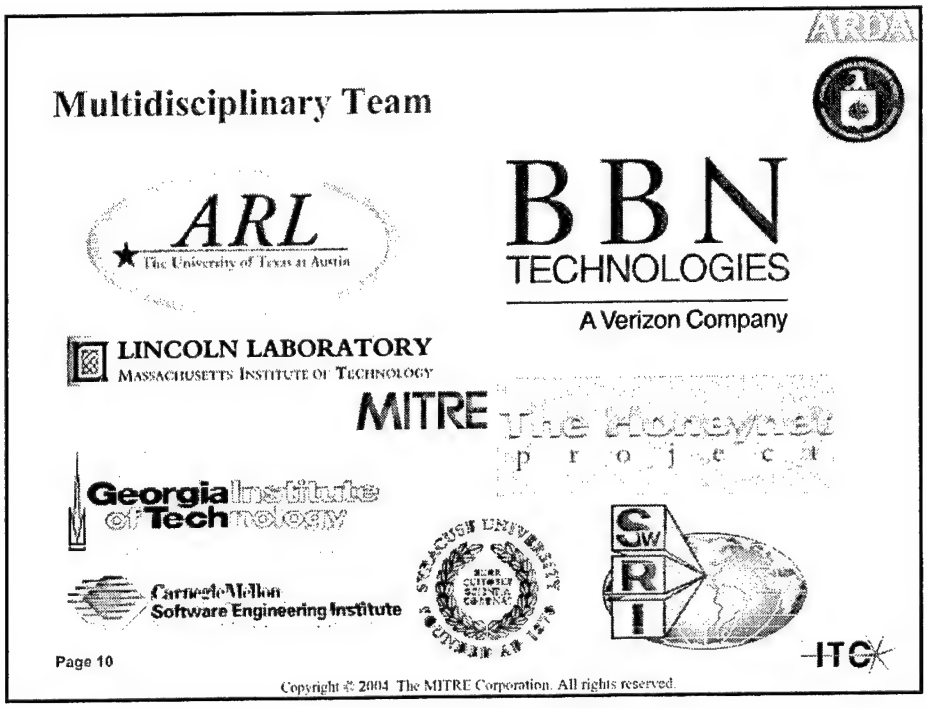
ARDA
Multidisciplinary Team
ARL
The University of Texas at Austin
BBN
TECHNOLOGIES
A Verizon Company
LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
MITRE
The Honeynet
project
Georgia Institute of Technology
Carnegie Mellon
Software Engineering Institute
SRI
ITC
Page 10
Copyright © 2004 The MITRE Corporation. All rights reserved.

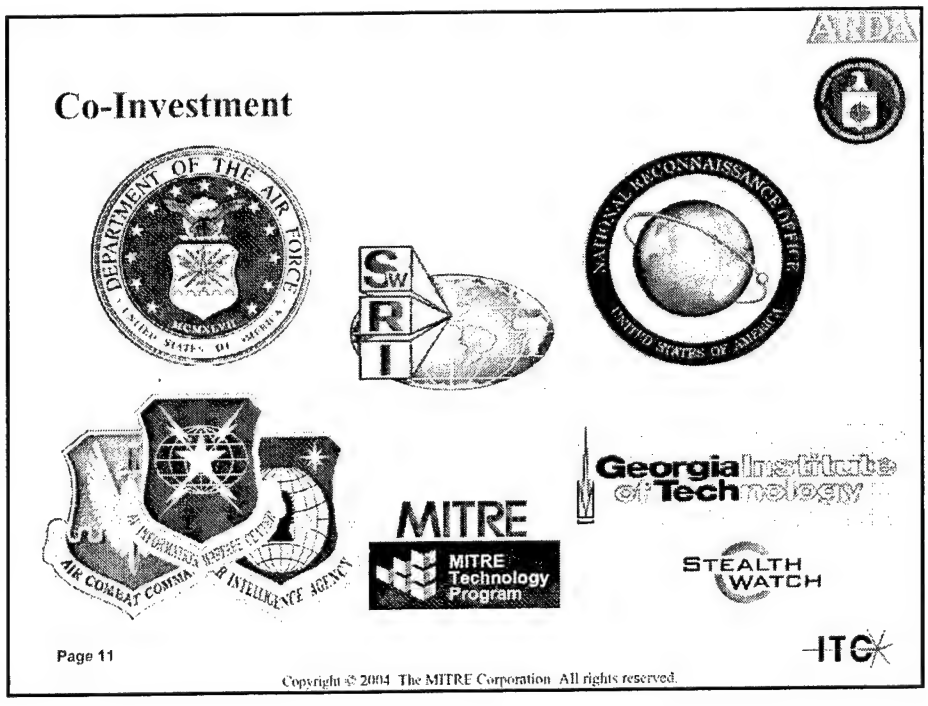
ARDA
Co-Investment
DEPARTMENT OF THE AIR FORCE
SRI
NATIONAL RECONNAISSANCE OFFICE
UNITED STATES OF AMERICA
MITRE
MITRE Technology Program
Georgia Institute of Technology
STEALTH WATCH
ITC
Page 11
Copyright © 2004 The MITRE Corporation. All rights reserved.

Workshop Participants

- Mr. Dick Brackney, ARDA
- Dr. Steve Chapin, Syracuse
- Dr. Brant Cheikes, MITRE
- Dr. John Copeland, Georgia Tech
- Mr. Mick Costa, MITRE
- Dr. Sandra Dykes, SwRI
- Mr. Thomas Eisenhut, SwRI
- Ms. Penny Lehtola, ARDA
- Mr. Jed Haile, Logan Group
- Mr. Tom Hetherington, ARL:UT
- Dr. Wenke Lee, Georgia Tech
- Mr. Scott Lewandowski, MIT LL
- Dr. Tom Longstaff, SEI of CMU
- Mrs. Paula MacDonald, MITRE
- Dr. Jack Marin, BBN
- Mrs. Sara Matzner, ARL:UT
- Dr. Mark Maybury, MITRE
- Mrs. Bev Nunan, MITRE
- Mr. Jeff Sebring, MITRE
- Mr. Conor Sibley, BBN
- Mr. Don Slife, consultant
- Mr. Lance Spitzner, Honeynet
- Mr. Jeffrey Picciotto, MITRE
- Mr. Richard Pietravalle, MITRE
- Mr. Brad Wood, BBN

And other graduate students and technical staff such as Christian Sarmoria and Cheol-min Hwang at Syracuse, undergraduate students George Chamales and Ryan Smith at ARL:UT; Bob Gaimari, Billy Garrison, and Laurie Damianos at MITRE

Page 12

Definitions

- Insider – **Anyone with access, privilege, or knowledge of information systems and services**
- Malicious insider (MI) – **motivated to intentionally adversely impact an organization's mission (e.g., deny, damage, degrade, destroy).**
- Observable - **Anything that can be detected with current tech.**
- Sensor – **Measures an observable (e.g., login, print, delete)**
  - **Sensor logs – recording of observables**
  - **Sensor stream – series of observables from one sensor**
- Indicator – **Identifiable event based on sensor output logs**
- Detect – **Determines event based on processing indicators**
- Report – **Indications and warnings of malicious insider behavior**
- Incident – **Related set of events**
- Fusion – **Processing multiple sensor outputs to provide a enhanced result (e.g., more abstract/concrete, higher confidence)**

Page 13

ITC

Hypotheses

- While some MIs can be detected using a single cyber observable, other MIs can be detected using a heterogeneous approach to indications and warning
  - Test using case analysis and experimentation
  - Increase the number of insiders detected
- Fusing information from heterogeneous information sources will allow us to formulate more accurate and timely indications and warning of insiders
  - Type of sensor (e.g., card reader, authentication, printer, telephone calls)
  - Level of IP stack (e.g., from network to application)
- Observables together with domain knowledge can help detect inappropriate behavior (e.g., need to know violations)
  - Domain model (e.g., user role, asset value to mission) helps distinguish need to know violations
  - Domain models need to be dynamic to distinguish the changing environment (user roles, asset value)

Group Hypotheses

- Structured Analysis
  - Real-time analysis of log data can classify some MI behavior as it occurs.
- Fusion
  - Accumulated cyber observables can be used to generate early indications of a MI.
- Honeynets
  - Instrumented targets can contribute to early MI identification.

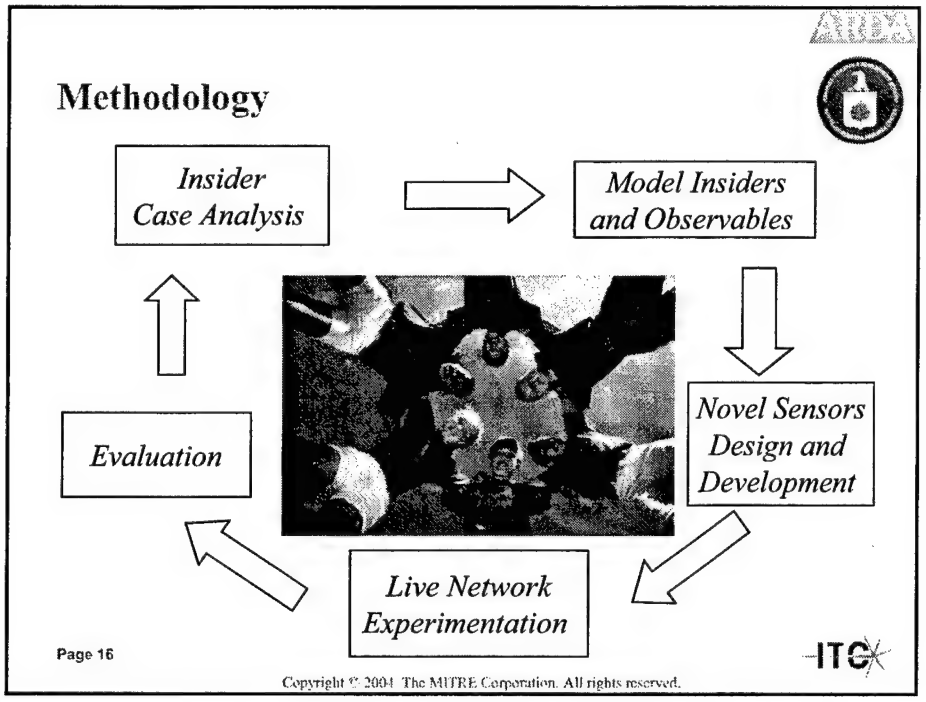
Methodology
Insider Case Analysis
Model Insiders and Observables
Novel Sensors Design and Development
Live Network Experimentation
Evaluation
Page 16
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

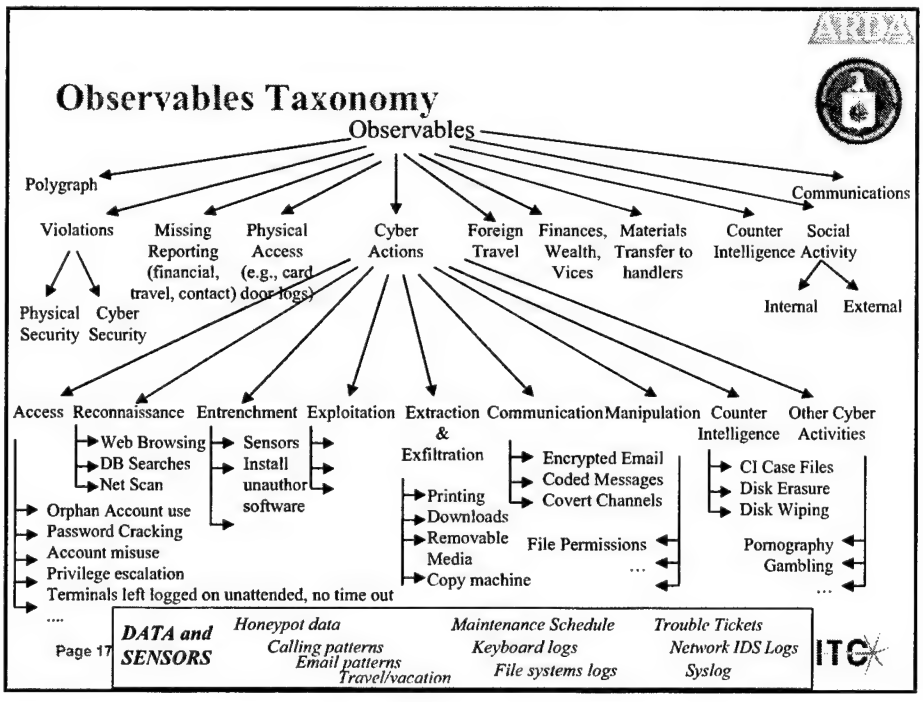
Observables Taxonomy
Observables
Polygraph
Communications
Violations
Missing Reporting (financial, travel, contact)
Physical Access (e.g., card door logs)
Cyber Actions
Foreign Travel
Finances, Wealth, Vices
Materials Transfer to handlers
Counter Intelligence
Social Activity
Physical Security
Cyber Security
Internal
External
Access
Reconnaissance
Entrenchment
Exploitation
Extraction & Exfiltration
Communication
Manipulation
Counter Intelligence
Other Cyber Activities
Web Browsing
DB Searches
Net Scan
Sensors
Install unauthor software
Encrypted Email
Coded Messages
Covert Channels
CI Case Files
Disk Erasure
Disk Wiping
Orphan Account use
Password Cracking
Account misuse
Privilege escalation
Terminals left logged on unattended, no time out
....
Printing
Downloads
Removable Media
Copy machine
File Permissions
...
Pornography
Gambling
...
DATA and SENSORS
Honeypot data
Calling patterns
Email patterns
Travel/vacation
Maintenance Schedule
Keyboard logs
File systems logs
Trouble Tickets
Network IDS Logs
Syslog
Page 17
ITC

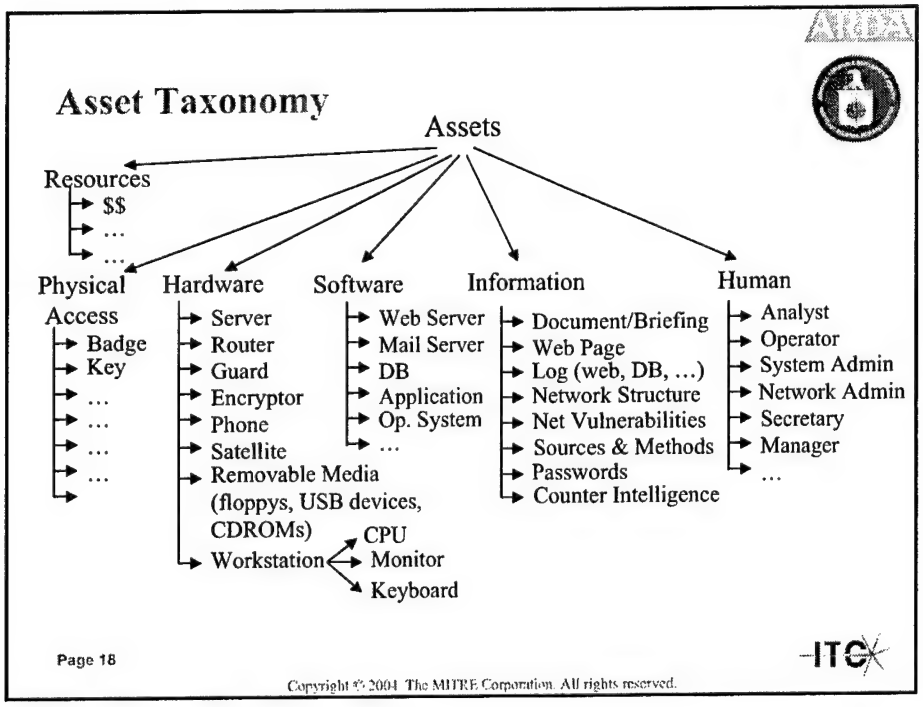
Asset Taxonomy
Assets
Resources
$$
…
…
Physical
Access
Badge
Key
…
…
…
…
Hardware
Server
Router
Guard
Encryptor
Phone
Satellite
Removable Media
(floppys, USB devices,
CDROMs)
Workstation
CPU
Monitor
Keyboard
Software
Web Server
Mail Server
DB
Application
Op. System
…
Information
Document/Briefing
Web Page
Log (web, DB, …)
Network Structure
Net Vulnerabilities
Sources & Methods
Passwords
Counter Intelligence
Human
Analyst
Operator
System Admin
Network Admin
Secretary
Manager
…
Page 18
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

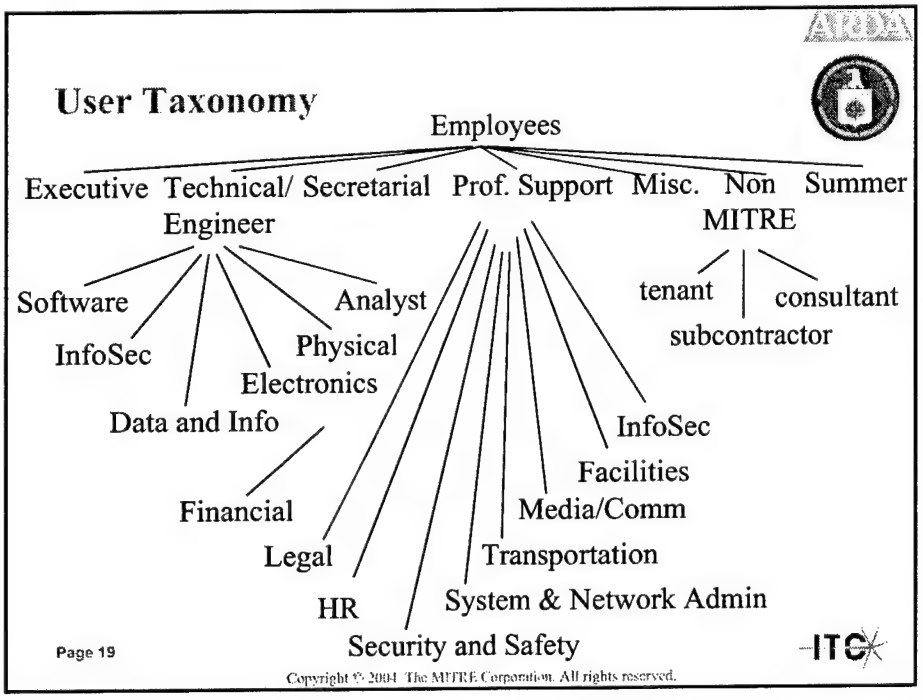
User Taxonomy
Employees
Executive
Technical/
Engineer
Secretarial
Prof. Support
Misc.
Non
MITRE
Summer
Software
InfoSec
Data and Info
Electronics
Physical
Analyst
Financial
Legal
HR
Security and Safety
System & Network Admin
Transportation
Media/Comm
Facilities
InfoSec
tenant
consultant
subcontractor
Page 19
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

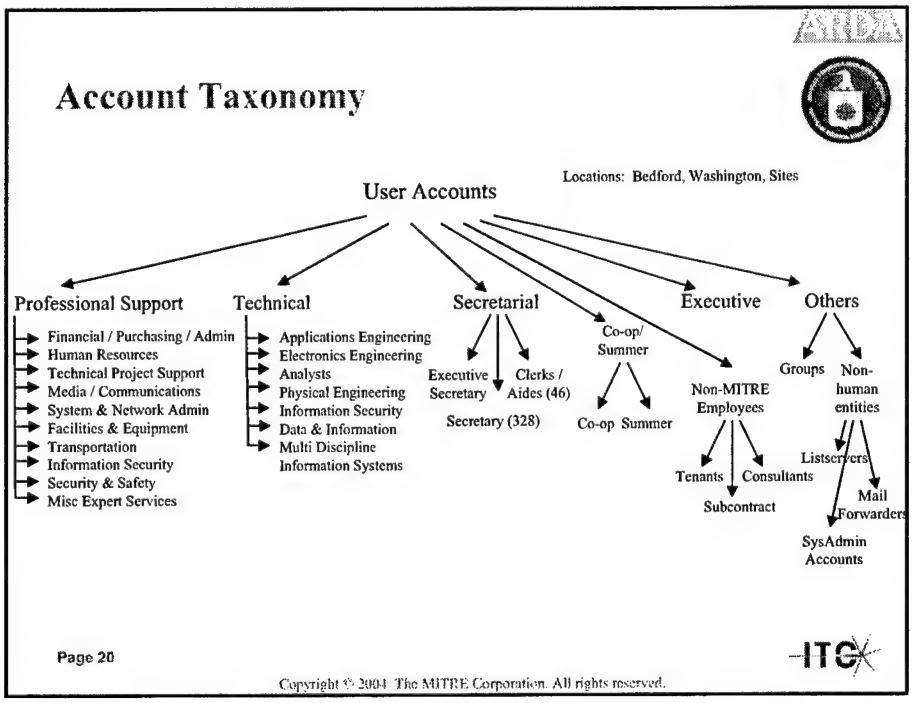
Account Taxonomy
Locations: Bedford, Washington, Sites
User Accounts
Professional Support
Financial / Purchasing / Admin
Human Resources
Technical Project Support
Media / Communications
System & Network Admin
Facilities & Equipment
Transportation
Information Security
Security & Safety
Misc Expert Services
Technical
Applications Engineering
Electronics Engineering
Analysts
Physical Engineering
Information Security
Data & Information
Multi Discipline
Information Systems
Secretarial
Executive Secretary
Clerks / Aides (46)
Secretary (328)
Co-op/ Summer
Co-op Summer
Executive
Non-MITRE Employees
Tenants
Consultants
Subcontract
Others
Groups
Non-human entities
Listservers
Mail Forwarders
SysAdmin Accounts
Page 20
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

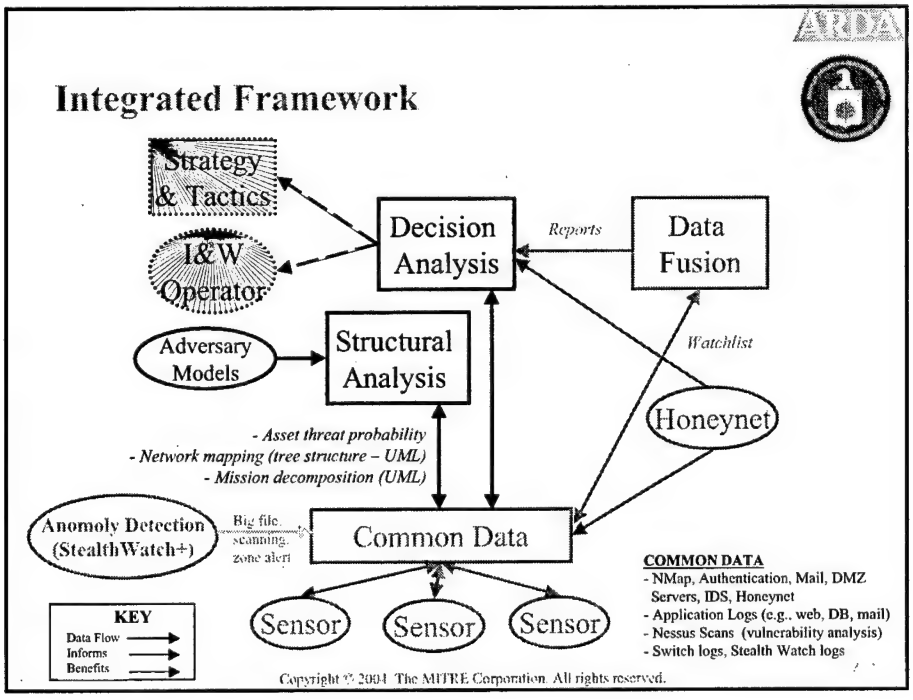
Integrated Framework
Strategy & Tactics
I&W Operator
Decision Analysis
Reports
Data Fusion
Adversary Models
Structural Analysis
Watchlist
Honeynet
- Asset threat probability
- Network mapping (tree structure – UML)
- Mission decomposition (UML)
Anomoly Detection (StealthWatch+)
Big file scanning, zone alert
Common Data
Sensor
Sensor
Sensor
KEY
Data Flow
Informs
Benefits
COMMON DATA
- NMap, Authentication, Mail, DMZ Servers, IDS, Honeynet
- Application Logs (e.g., web, DB, mail)
- Nessus Scans (vulnerability analysis)
- Switch logs, Stealth Watch logs
Copyright © 2004 The MITRE Corporation. All rights reserved.

Malicious Insider Data Set

- Real network - MITRE's DMZ
  - A separate network for experimentation and sponsor community support established outside of the MITRE internal network
  - 300 – 400 hosts
  - Various services: Web, news, email, database, ...
  - Data sources on network for use in scenarios
  - Deploy additional sensors
- Three of 75 users active during period acted as malicious insiders based on historical and project scenarios of insider behavior

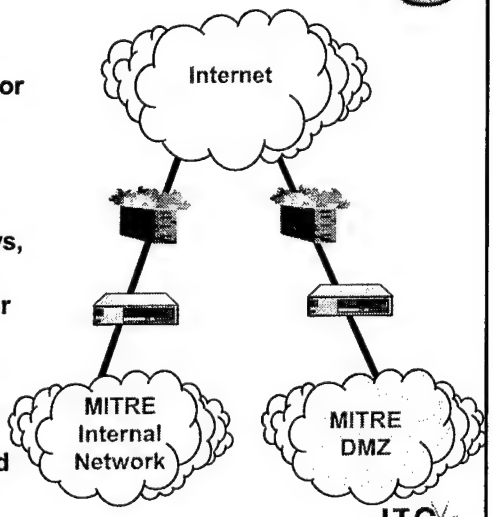

ITC

Testbed Network

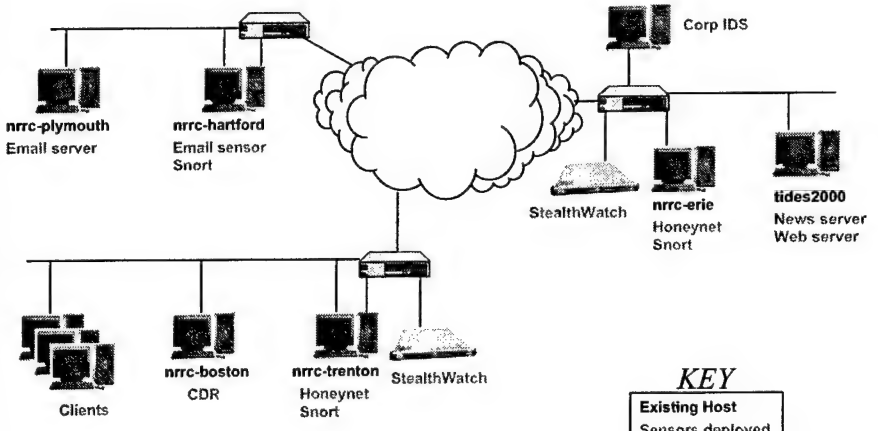

Page 23

ITC

Our Insider Knowledge and Focus

STEALTH	Occurred	Not yet Occurred
Detectable	Robert Philip Hanssen	MI who attacks the network
Hard to Detect	~~Ana Belen Montes~~	~~Non-cyber component~~
Not yet Detectable		

X - Unaddressed by workshop
X – Unobservable in cyberspace

Focus: "Indications and warnings not conviction and sentencing"

Page 24

ITC

Insider Scenarios

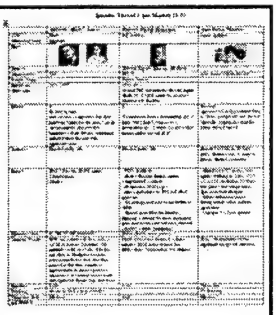

- Three scenarios:
  - Aggregate Historical Insider
    - "Pal"
  - Projected Insiders
    - TIDES Admin
    - "Jack"

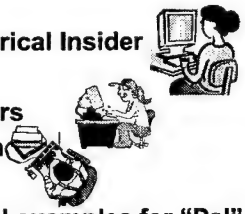

- Drew upon historical examples for "Pal"
  - Intelligence analyst
- TIDES Admin and "Jack" developed their scenarios
  - Needed to be consistent with prior activity on systems
    - An application administrator
    - A system administrator
  - More realistic ("red teaming")

Page 25

ITC

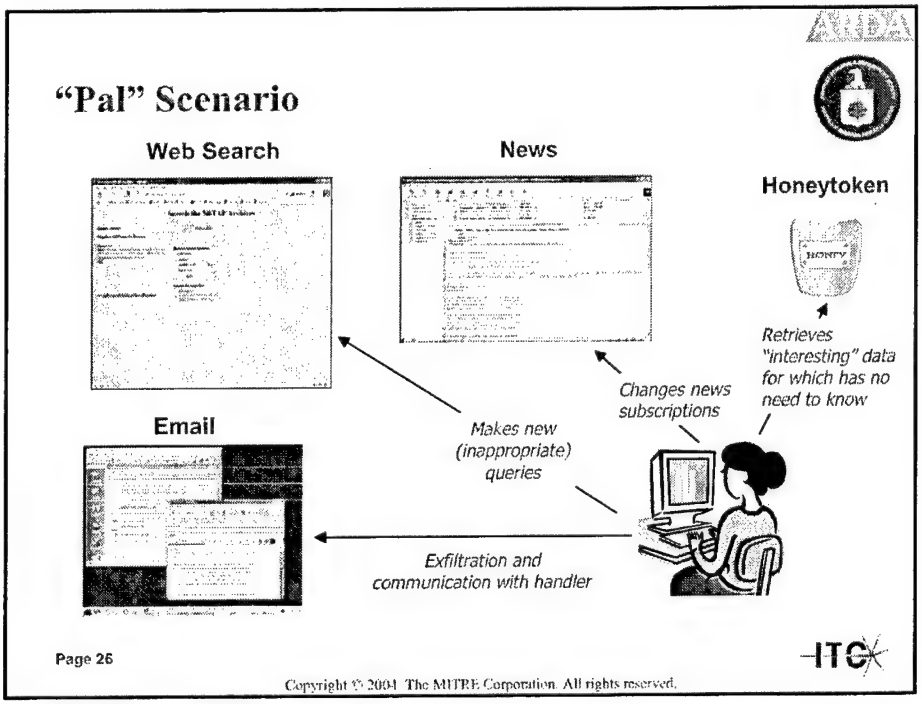
"Pal" Scenario
Web Search
News
Honeytoken
HONEY
Retrieves "interesting" data for which has no need to know
Changes news subscriptions
Email
Makes new (inappropriate) queries
Exfiltration and communication with handler
Page 26
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

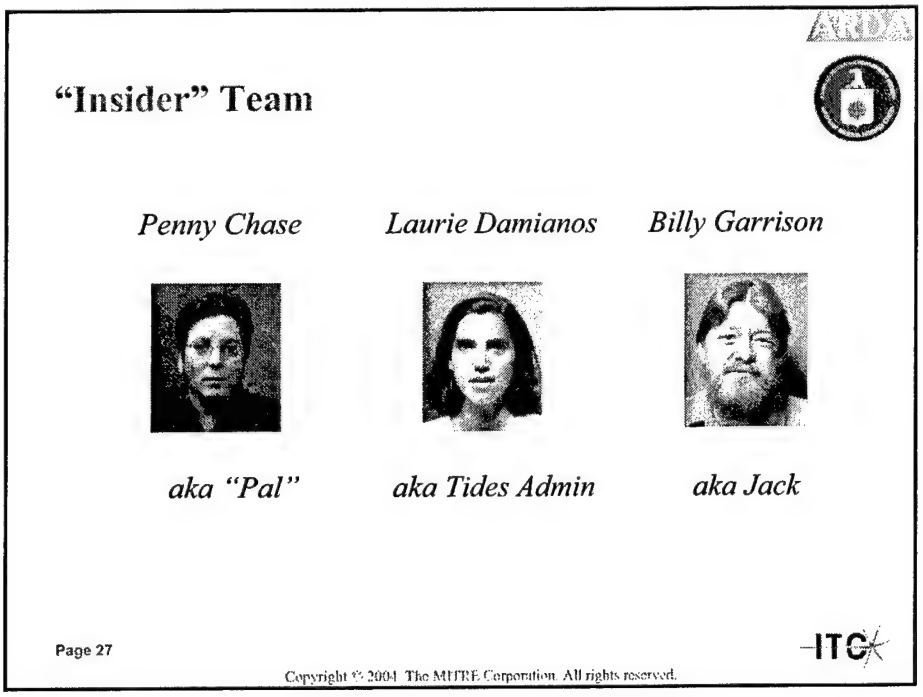
"Insider" Team
Penny Chase
Laurie Damianos
Billy Garrison
aka "Pal"
aka Tides Admin
aka Jack
Page 27
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

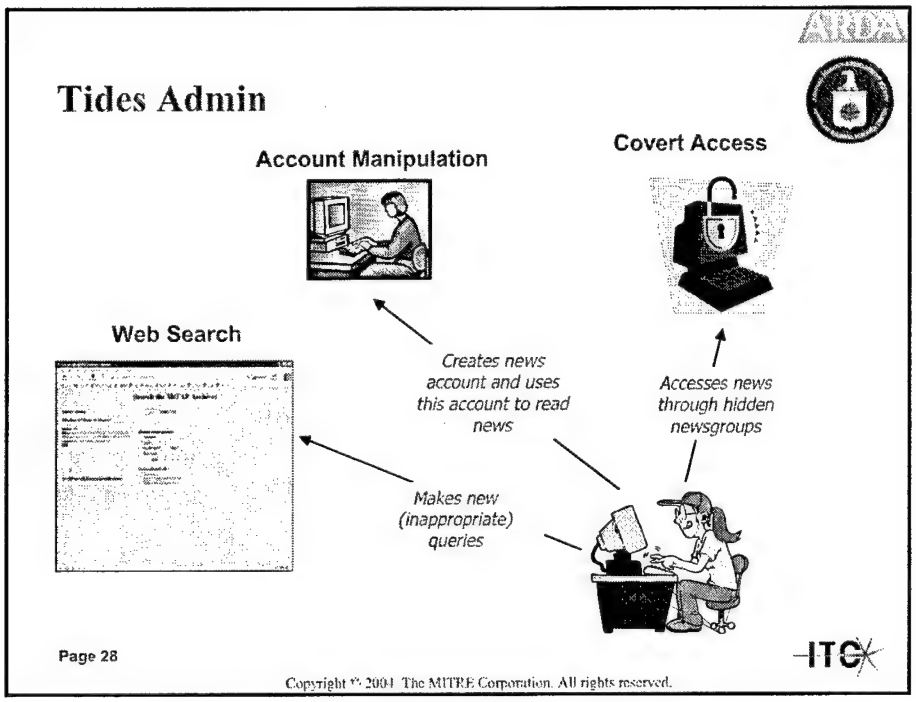
Tides Admin
Account Manipulation
Covert Access
Web Search
Creates news account and uses this account to read news
Accesses news through hidden newsgroups
Makes new (inappropriate) queries
Page 28
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

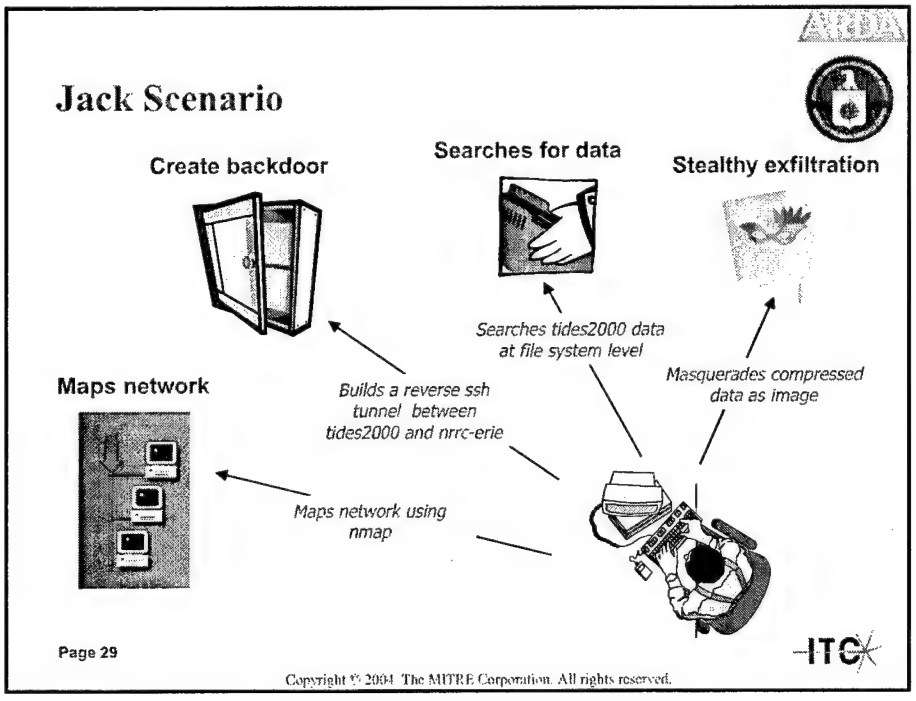
Jack Scenario
Create backdoor
Searches for data
Stealthy exfiltration
Searches tides2000 data at file system level
Masquerades compressed data as image
Maps network
Builds a reverse ssh tunnel between tides2000 and nrrc-erie
Maps network using nmap
Page 29
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

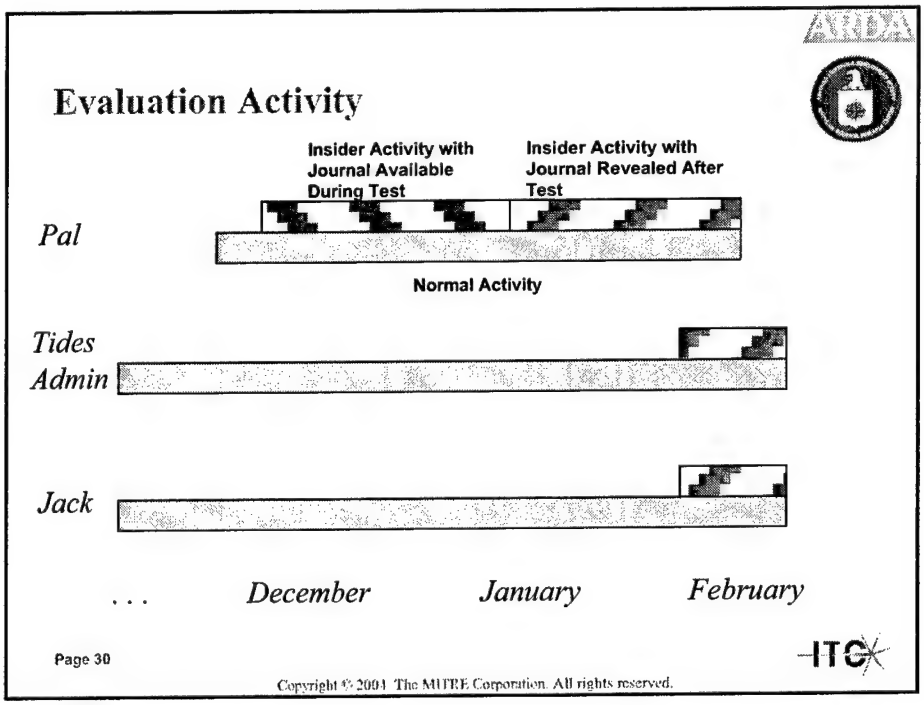
Evaluation Activity
Insider Activity with Journal Available During Test
Insider Activity with Journal Revealed After Test
Pal
Normal Activity
Tides Admin
Jack
. . .
December
January
February
Page 30
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

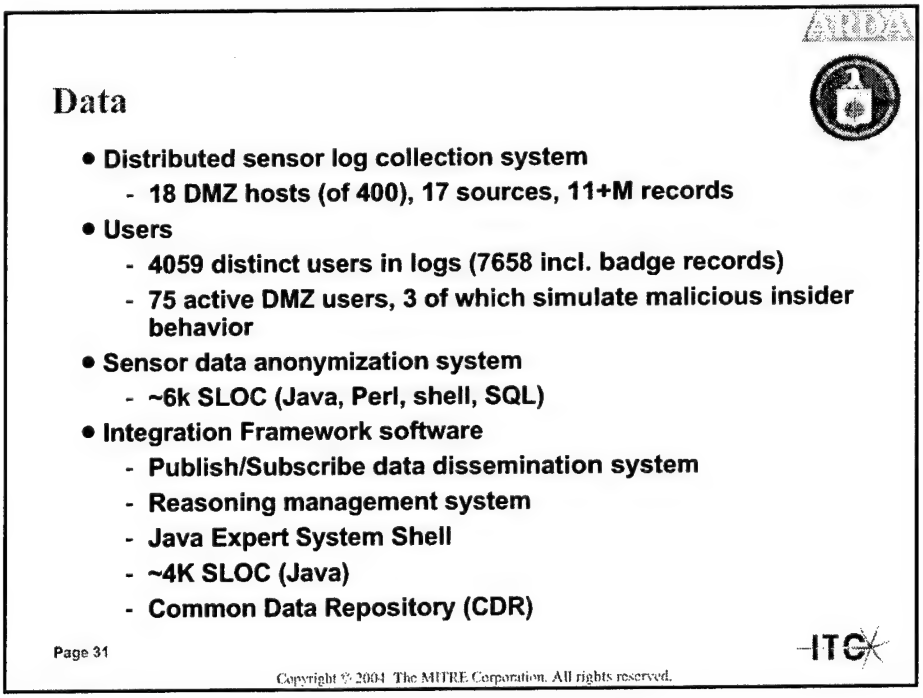
Data
• Distributed sensor log collection system
- 18 DMZ hosts (of 400), 17 sources, 11+M records
• Users
- 4059 distinct users in logs (7658 incl. badge records)
- 75 active DMZ users, 3 of which simulate malicious insider behavior
• Sensor data anonymization system
- ~6k SLOC (Java, Perl, shell, SQL)
• Integration Framework software
- Publish/Subscribe data dissemination system
- Reasoning management system
- Java Expert System Shell
- ~4K SLOC (Java)
- Common Data Repository (CDR)
Page 31
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

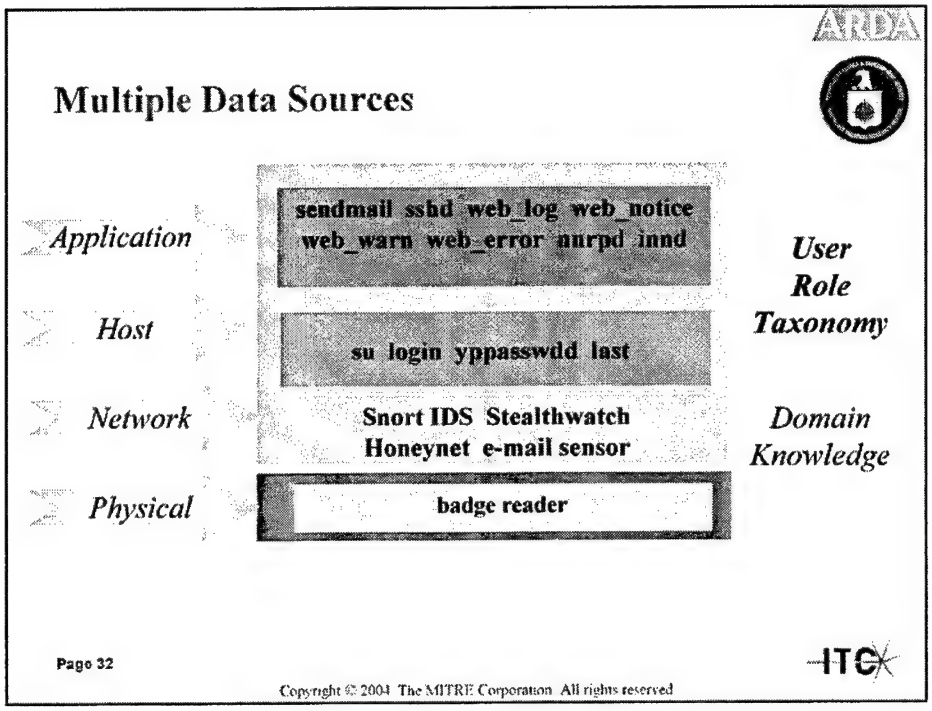
ARDA
Multiple Data Sources
Application
Host
Network
Physical
sendmail sshd web_log web_notice
web_warn web_error nurpd innd
su login yppasswdd last
Snort IDS Stealthwatch
Honeynet e-mail sensor
badge reader
User
Role
Taxonomy
Domain
Knowledge
Page 32
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved

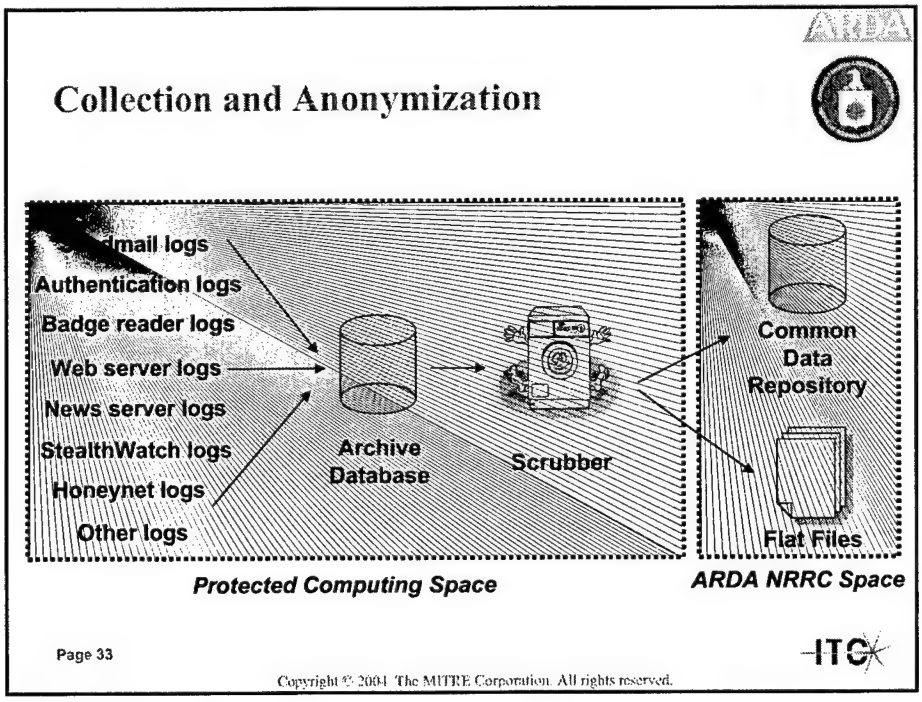
ARDA
Collection and Anonymization
mail logs
Authentication logs
Badge reader logs
Web server logs
News server logs
StealthWatch logs
Honeynet logs
Other logs
Archive
Database
Scrubber
Common
Data
Repository
Flat Files
Protected Computing Space
ARDA NRRC Space
Page 33
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

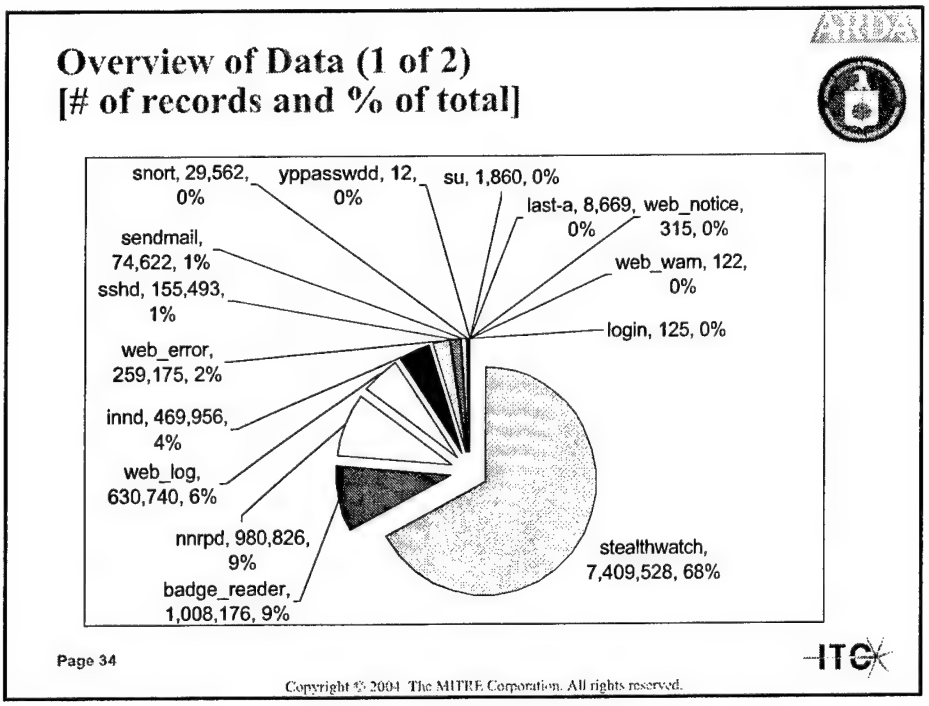
Overview of Data (1 of 2)
[# of records and % of total]
snort, 29,562, 0%
yppasswdd, 12, 0%
su, 1,860, 0%
last-a, 8,669, 0%
web_notice, 315, 0%
sendmail, 74,622, 1%
web_warn, 122, 0%
sshd, 155,493, 1%
login, 125, 0%
web_error, 259,175, 2%
innd, 469,956, 4%
web_log, 630,740, 6%
nnrpd, 980,826, 9%
badge_reader, 1,008,176, 9%
stealthwatch, 7,409,528, 68%
Page 34
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

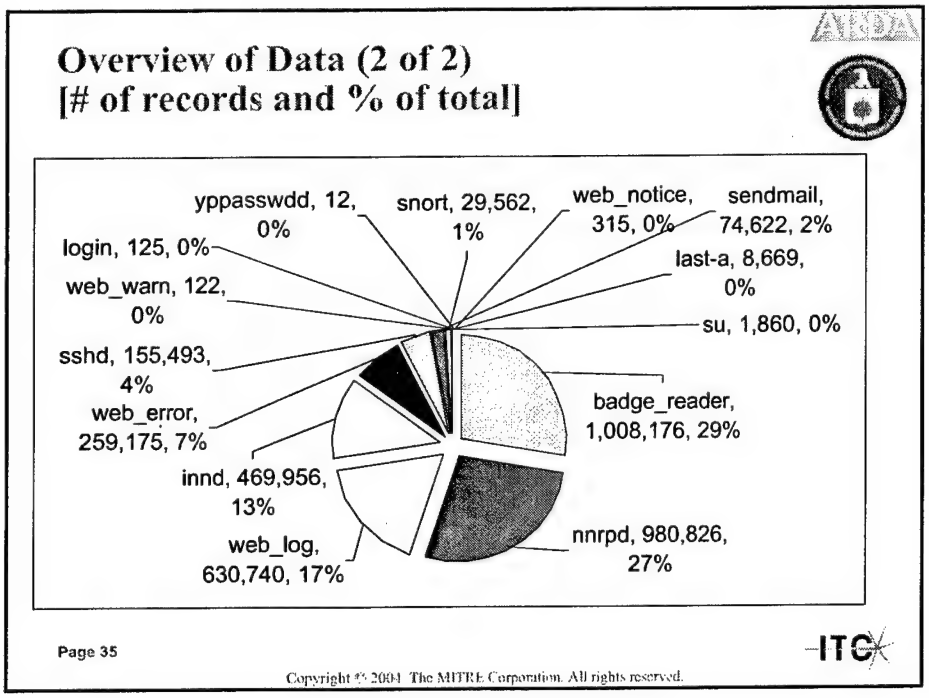
Overview of Data (2 of 2)
[# of records and % of total]
yppasswdd, 12, 0%
snort, 29,562, 1%
web_notice, 315, 0%
sendmail, 74,622, 2%
login, 125, 0%
last-a, 8,669, 0%
web_warn, 122, 0%
su, 1,860, 0%
sshd, 155,493, 4%
badge_reader, 1,008,176, 29%
web_error, 259,175, 7%
innd, 469,956, 13%
nnrpd, 980,826, 27%
web_log, 630,740, 17%
Page 35
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

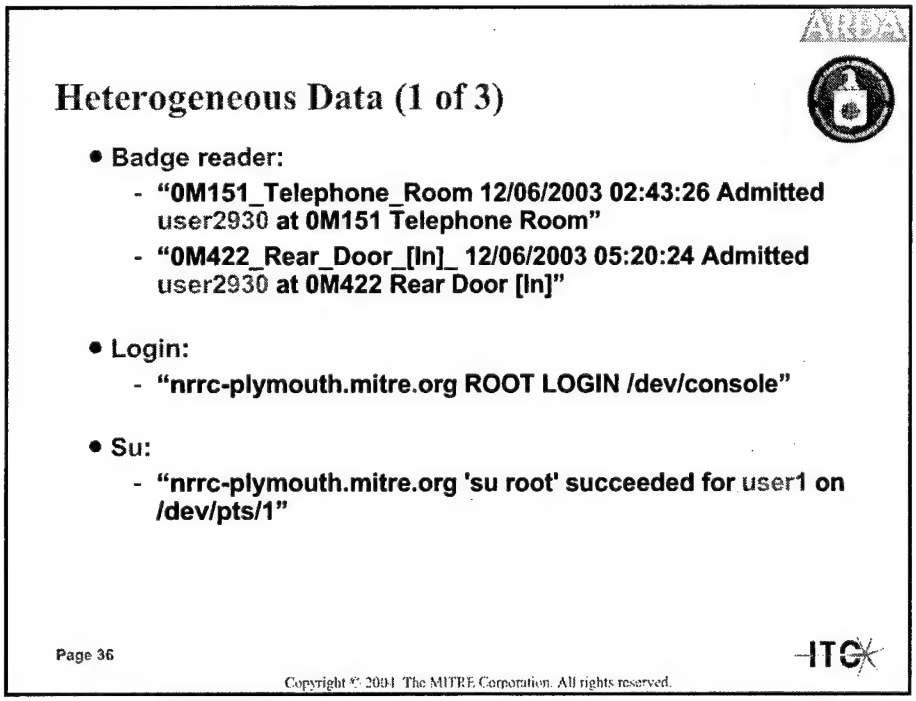
ARDA
Heterogeneous Data (1 of 3)
• Badge reader:
- "0M151_Telephone_Room 12/06/2003 02:43:26 Admitted user2930 at 0M151 Telephone Room"
- "0M422_Rear_Door_[In]_ 12/06/2003 05:20:24 Admitted user2930 at 0M422 Rear Door [In]"
• Login:
- "nrrc-plymouth.mitre.org ROOT LOGIN /dev/console"
• Su:
- "nrrc-plymouth.mitre.org 'su root' succeeded for user1 on /dev/pts/1"
Page 36
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

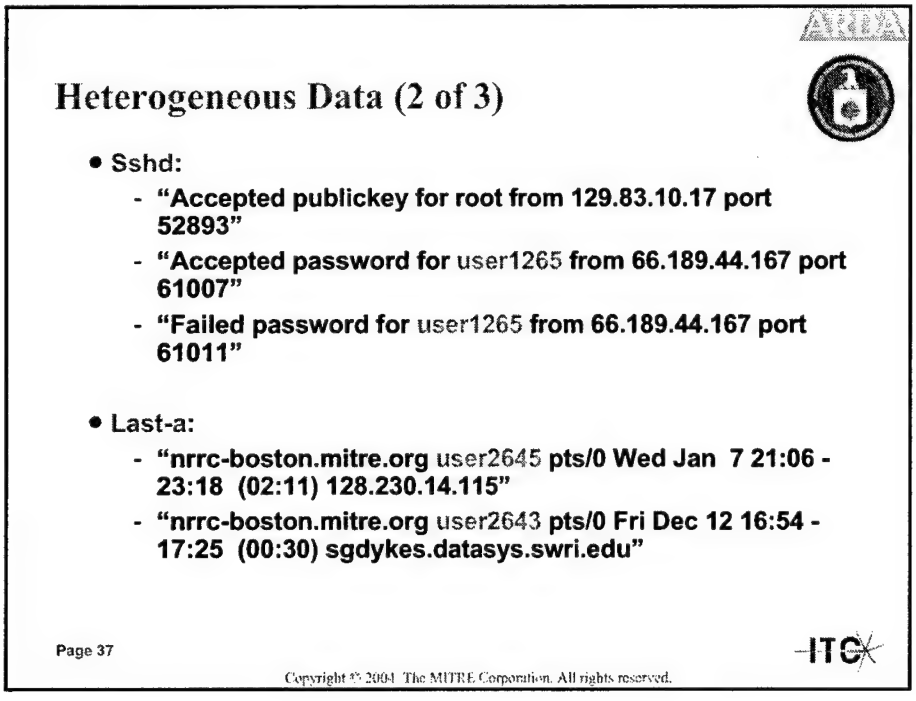
ARDA
Heterogeneous Data (2 of 3)
• Sshd:
- "Accepted publickey for root from 129.83.10.17 port 52893"
- "Accepted password for user1265 from 66.189.44.167 port 61007"
- "Failed password for user1265 from 66.189.44.167 port 61011"
• Last-a:
- "nrrc-boston.mitre.org user2645 pts/0 Wed Jan 7 21:06 - 23:18 (02:11) 128.230.14.115"
- "nrrc-boston.mitre.org user2643 pts/0 Fri Dec 12 16:54 - 17:25 (00:30) sgdykes.datasys.swri.edu"
Page 37
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

Heterogeneous Data (3 of 3)

- Web_log:
  - **"GET /cvw/licenses/source/license.html HTTP/1.0"**
  - **"GET /basilix.php3?request_id[DUMMY]=../../../../etc/passwd &RequestID=DUMMY&username=user2311&password=xxxxx HTTP/1.1"**
- Web_error:
  - **"Invalid method in request get /scripts/..."**
  - **"File does not exist: /tides_1/.../etc/passwd"**
- Sendmail:
  - **"cvw.mitre.org 14436 i0J507Lb014436: from=<user10368@digito.com>, size=2789, class=0, nrcpts=0, proto=ESMTP, daemon=MTA, relay=smtp-bedford-x.mitre.org [192.160.51.76]"**
  - **"cvw.mitre.org 14645 i0J7ErLb014644: to=user8, ctladdr=<user9@cvw.mitre.org> (1/0), delay=00:00:00, xdelay=00:00:00, mailer=*file*, pri=41013, dsn=2.0.0, stat=Sent"**

Page 38

Integration Framework

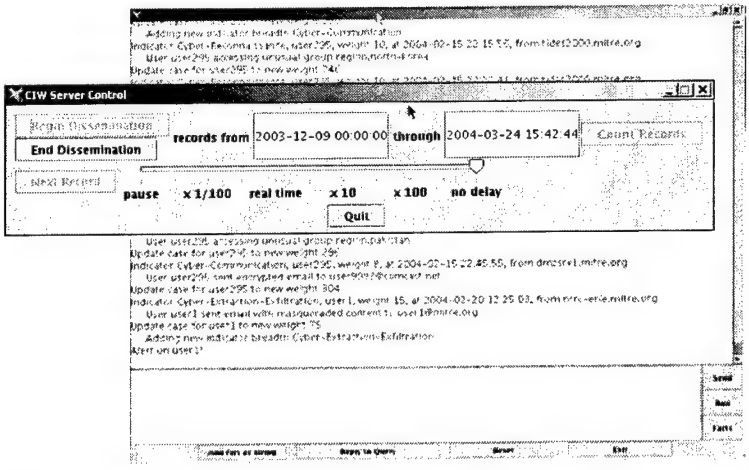

Page 39

Who is Suspicious?

user1265
user8859
user301
user2304
user2645
user324
user2306
user215
user2649
user268
user2644
user287
user319
user322
user252
user2647
user11838
user317
user2648

Page 40

I&W Approaches

- **StealthWatch**
  - Multilevel Monitoring – packet level (John Copeland, Georgia Tech), system level, and application level.
- **Honeynets**
  - Distributed honeynets to acquire attacker properties, pre-attack intensions, and potential attack strategies (Lance Spitzner/Jed Haile, Honeynet
- **Structured Analysis Group (SAG)**
  - Novel functional model related to attack graphs which will map pre-attack indicators to potential attacks (Brad Wood/Jack Marin, BBN; Steve Chapin, Syracuse)
- **Data Fusion**
  - Automatic fusion of traditional and novel indicators (Sara Matzner, U Texas; Sandy Dykes, SwRI)

Page 41

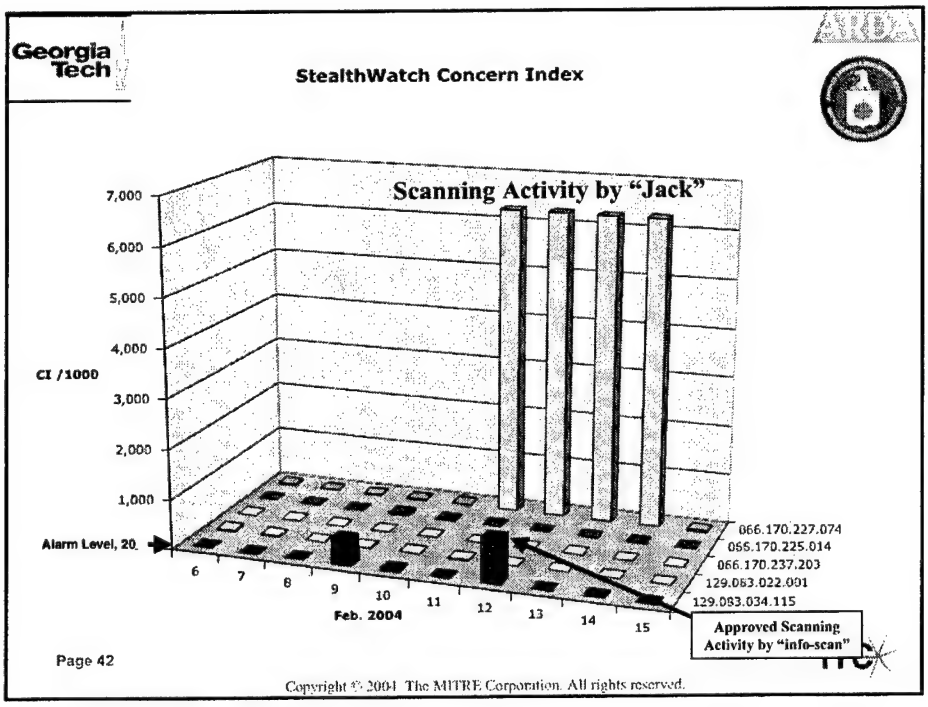
Georgia Tech
StealthWatch Concern Index
Scanning Activity by "Jack"
7,000
6,000
5,000
4,000
CI /1000
3,000
2,000
1,000
Alarm Level, 20
6
7
8
9
10
11
12
13
14
15
Feb. 2004
066.170.227.074
066.170.225.014
066.170.237.203
129.083.022.001
129.083.034.115
Approved Scanning Activity by "info-scan"
Page 42
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

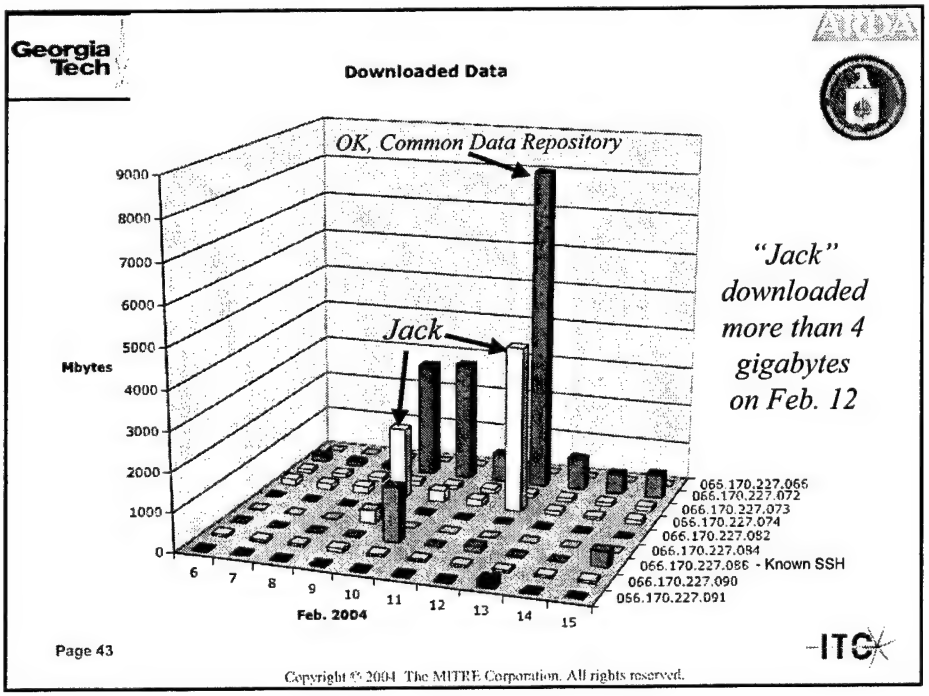
Georgia Tech
Downloaded Data
OK, Common Data Repository
9000
8000
7000
6000
Jack
5000
Mbytes
4000
3000
2000
1000
0
6
7
8
9
10
11
12
13
14
15
Feb. 2004
066.170.227.066
066.170.227.072
066.170.227.073
066.170.227.074
066.170.227.082
066.170.227.084
066.170.227.086 - Known SSH
066.170.227.090
066.170.227.091
"Jack" downloaded more than 4 gigabytes on Feb. 12
Page 43
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

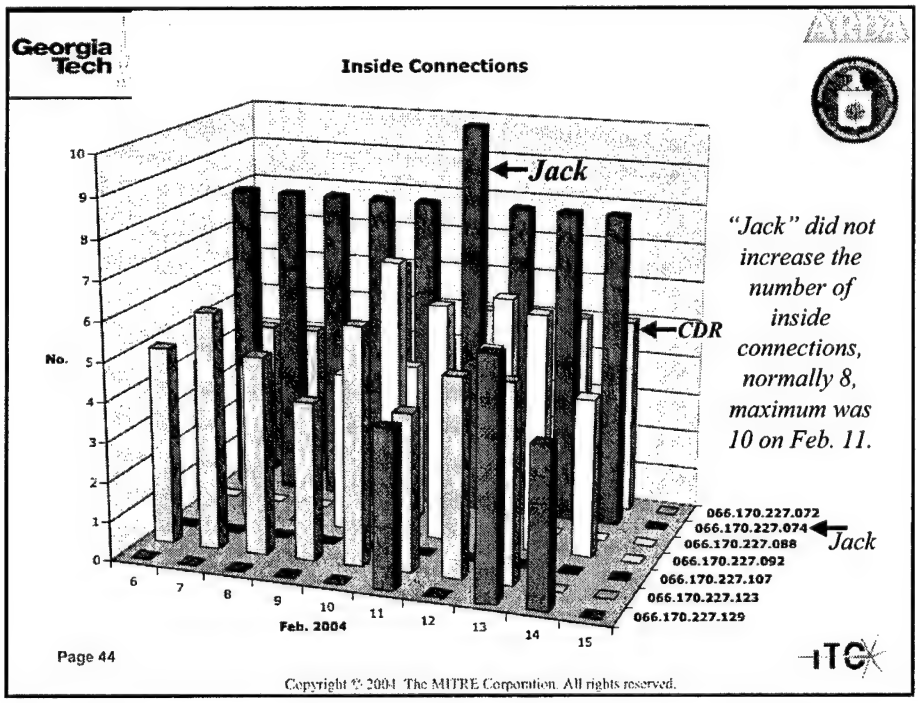
Georgia Tech
Inside Connections
Jack
CDR
No.
10
9
8
7
6
5
4
3
2
1
0
6
7
8
9
10
11
12
13
14
15
Feb. 2004
066.170.227.072
066.170.227.074
066.170.227.088
066.170.227.092
066.170.227.107
066.170.227.123
066.170.227.129
Jack
"Jack" did not increase the number of inside connections, normally 8, maximum was 10 on Feb. 11.
Page 44
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

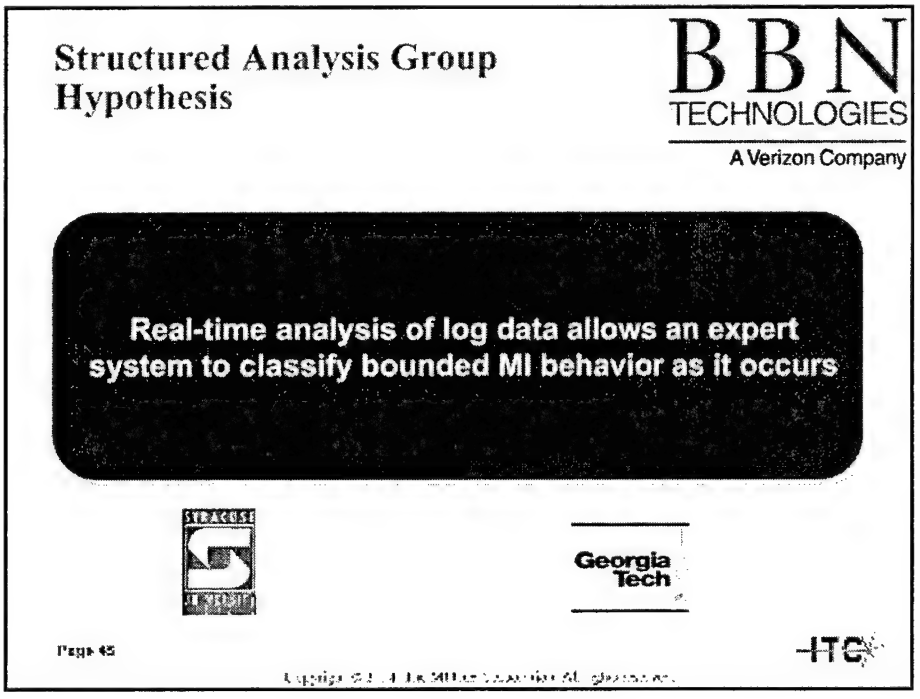
Structured Analysis Group
Hypothesis
BBN
TECHNOLOGIES
A Verizon Company
Real-time analysis of log data allows an expert system to classify bounded MI behavior as it occurs
Georgia Tech
ITC

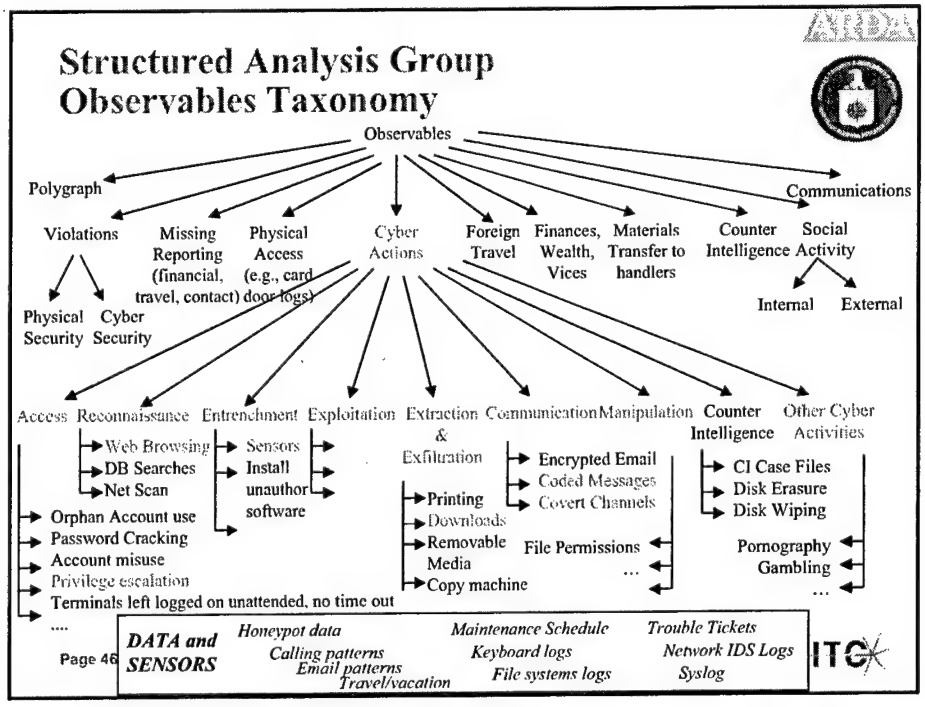

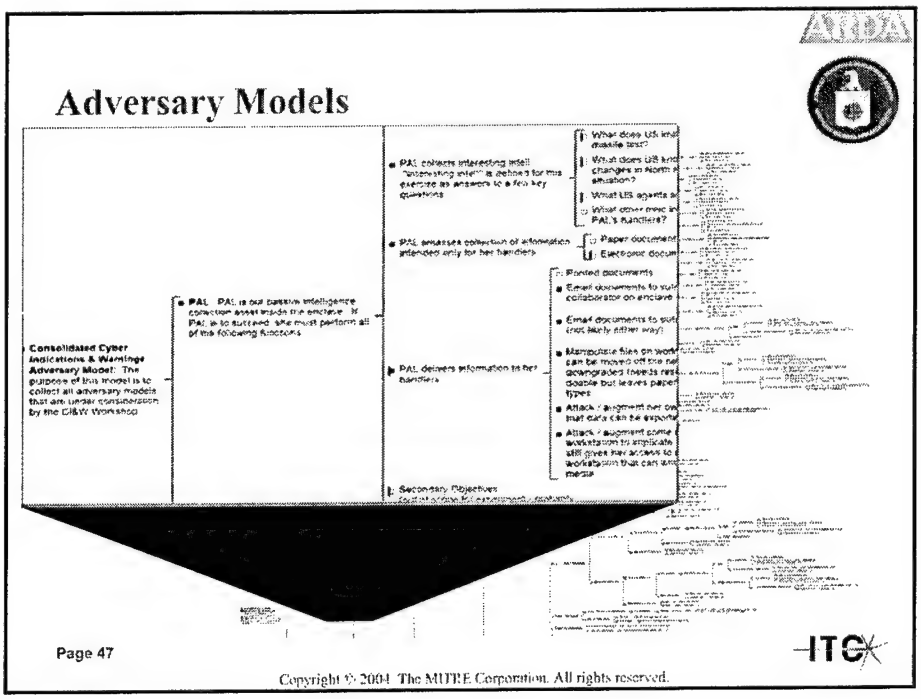

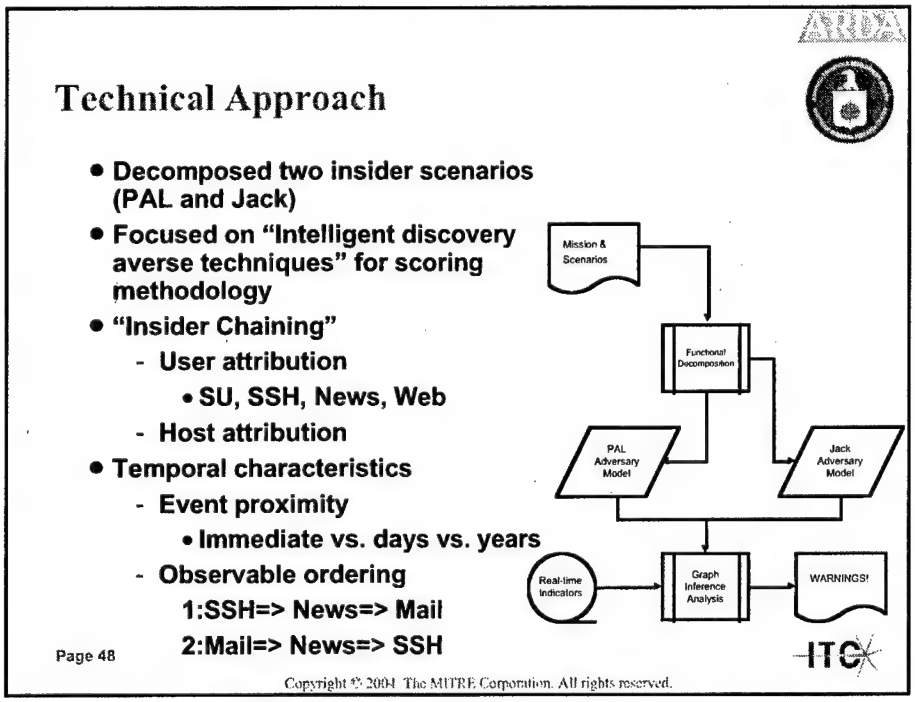

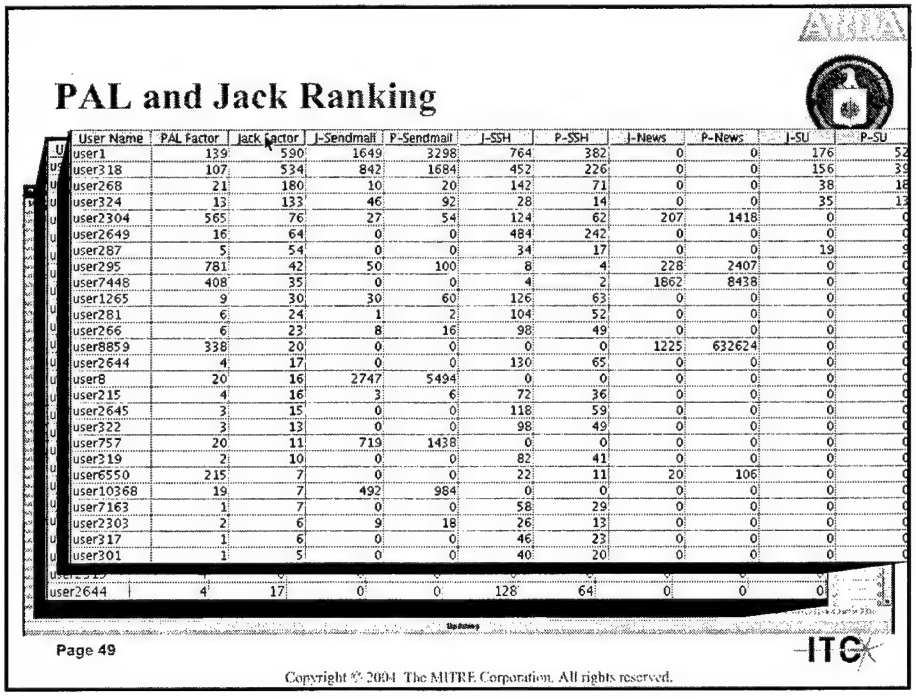

User Name	PAL Factor	Jack Factor	J-Sendmail	P-Sendmail	J-SSH	P-SSH	J-News	P-News	J-SU	P-SU
user1	139	590	1649	3298	764	382	0	0	176	52
user318	107	534	842	1684	452	226	0	0	156	39
user268	21	180	10	20	142	71	0	0	38	18
user324	13	133	46	92	28	14	0	0	35	13
user2304	565	76	27	54	124	62	207	1418	0	0
user2649	16	64	0	0	484	242	0	0	0	0
user287	5	54	0	0	34	17	0	0	19	9
user295	781	42	50	100	8	4	228	2407	0	0
user7448	408	35	0	0	4	2	1862	8438	0	0
user1265	9	30	30	60	126	63	0	0	0	0
user281	6	24	1	2	104	52	0	0	0	0
user266	6	23	8	16	98	49	0	0	0	0
user8859	338	20	0	0	0	0	1225	632624	0	0
user2644	4	17	0	0	130	65	0	0	0	0
user8	20	16	2747	5494	0	0	0	0	0	0
user215	4	16	3	6	72	36	0	0	0	0
user2645	3	15	0	0	118	59	0	0	0	0
user322	3	13	0	0	98	49	0	0	0	0
user757	20	11	719	1438	0	0	0	0	0	0
user319	2	10	0	0	82	41	0	0	0	0
user6550	215	7	0	0	22	11	20	106	0	0
user10368	19	7	492	984	0	0	0	0	0	0
user7163	1	7	0	0	58	29	0	0	0	0
user2303	2	6	9	18	26	13	0	0	0	0
user317	1	6	0	0	46	23	0	0	0	0
user301	1	5	0	0	40	20	0	0	0	0
user2644	4	17	0	0	128	64	0	0	0	0

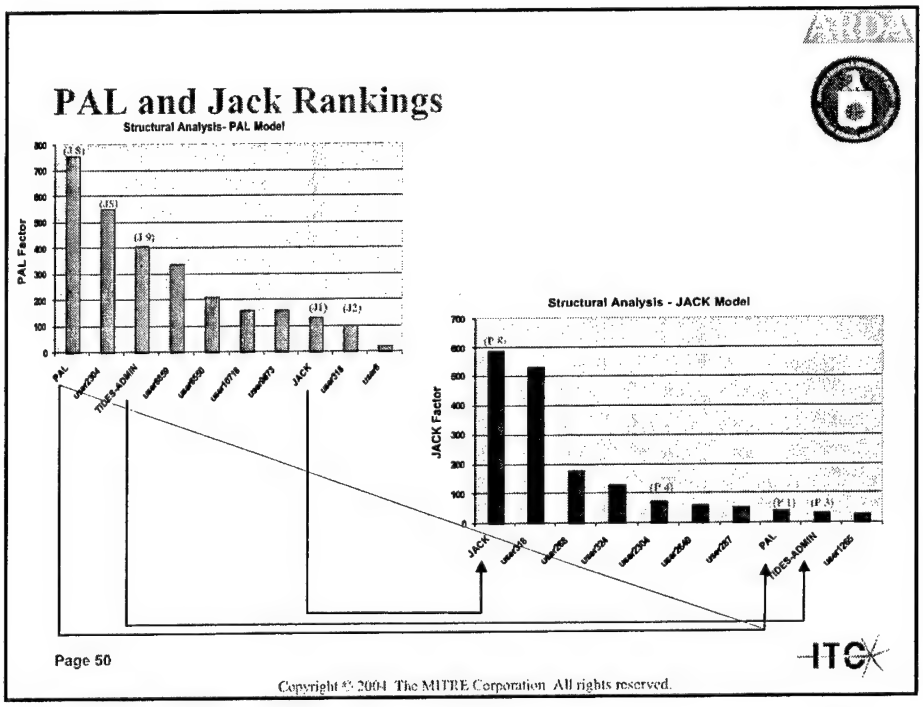
PAL and Jack Rankings
Structural Analysis- PAL Model
PAL Factor
Structural Analysis - JACK Model
JACK Factor
Page 50
ITC
Copyright © 2004 The MITRE Corporation All rights reserved.

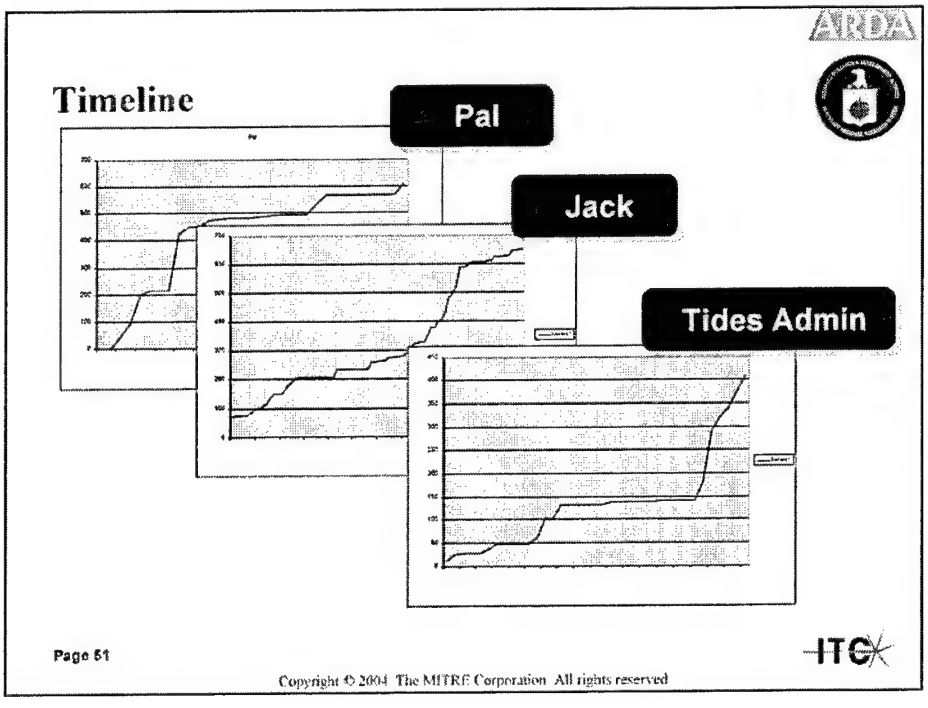
Timeline
Pal
Jack
Tides Admin
Page 51
ITC
Copyright © 2004 The MITRE Corporation All rights reserved

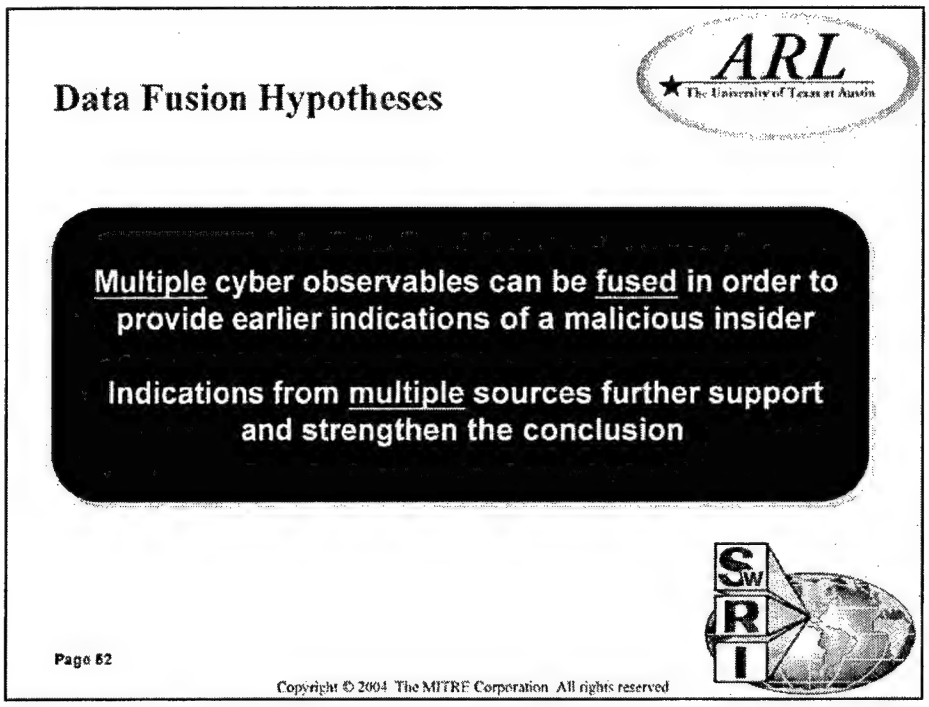
Data Fusion Hypotheses
ARL
The University of Texas at Austin
Multiple cyber observables can be fused in order to provide earlier indications of a malicious insider
Indications from multiple sources further support and strengthen the conclusion
SwRI
Page 52
Copyright © 2004 The MITRE Corporation All rights reserved

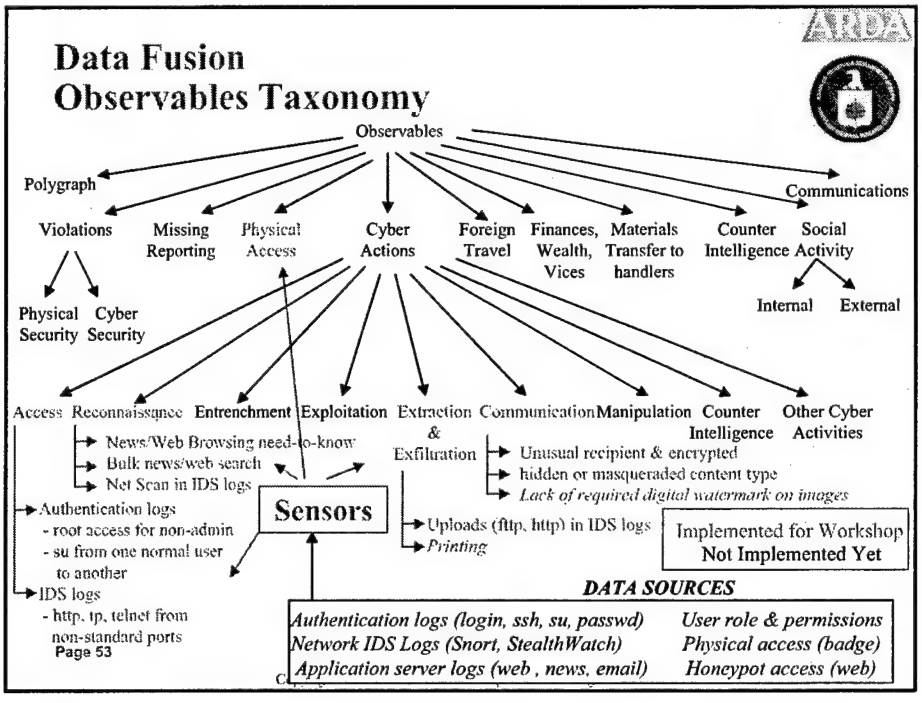
ARDA
Data Fusion
Observables Taxonomy
Observables
Polygraph
Communications
Violations
Missing Reporting
Physical Access
Cyber Actions
Foreign Travel
Finances, Wealth, Vices
Materials Transfer to handlers
Counter Intelligence
Social Activity
Physical Security
Cyber Security
Internal
External
Access
Reconnaissance
Entrenchment
Exploitation
Extraction & Exfiltration
Communication
Manipulation
Counter Intelligence
Other Cyber Activities
News/Web Browsing need-to-know
Bulk news/web search
Net Scan in IDS logs
Unusual recipient & encrypted
hidden or masqueraded content type
Lack of required digital watermark on images
Authentication logs
- root access for non-admin
- su from one normal user to another
IDS logs
- http, tp, telnet from non-standard ports
Sensors
Uploads (ftp, http) in IDS logs
Printing
Implemented for Workshop
Not Implemented Yet
DATA SOURCES
Authentication logs (login, ssh, su, passwd)
Network IDS Logs (Snort, StealthWatch)
Application server logs (web , news, email)
User role & permissions
Physical access (badge)
Honeypot access (web)
Page 53

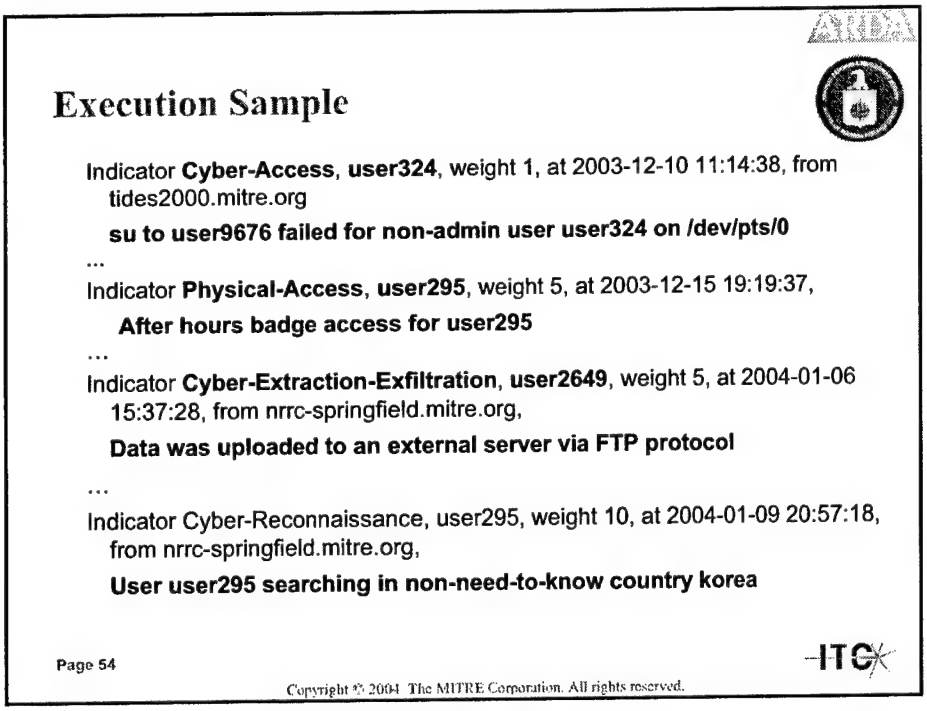
Execution Sample
Indicator Cyber-Access, user324, weight 1, at 2003-12-10 11:14:38, from tides2000.mitre.org
su to user9676 failed for non-admin user user324 on /dev/pts/0
...
Indicator Physical-Access, user295, weight 5, at 2003-12-15 19:19:37,
After hours badge access for user295
...
Indicator Cyber-Extraction-Exfiltration, user2649, weight 5, at 2004-01-06 15:37:28, from nrrc-springfield.mitre.org,
Data was uploaded to an external server via FTP protocol
...
Indicator Cyber-Reconnaissance, user295, weight 10, at 2004-01-09 20:57:18, from nrrc-springfield.mitre.org,
User user295 searching in non-need-to-know country korea
Page 54
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

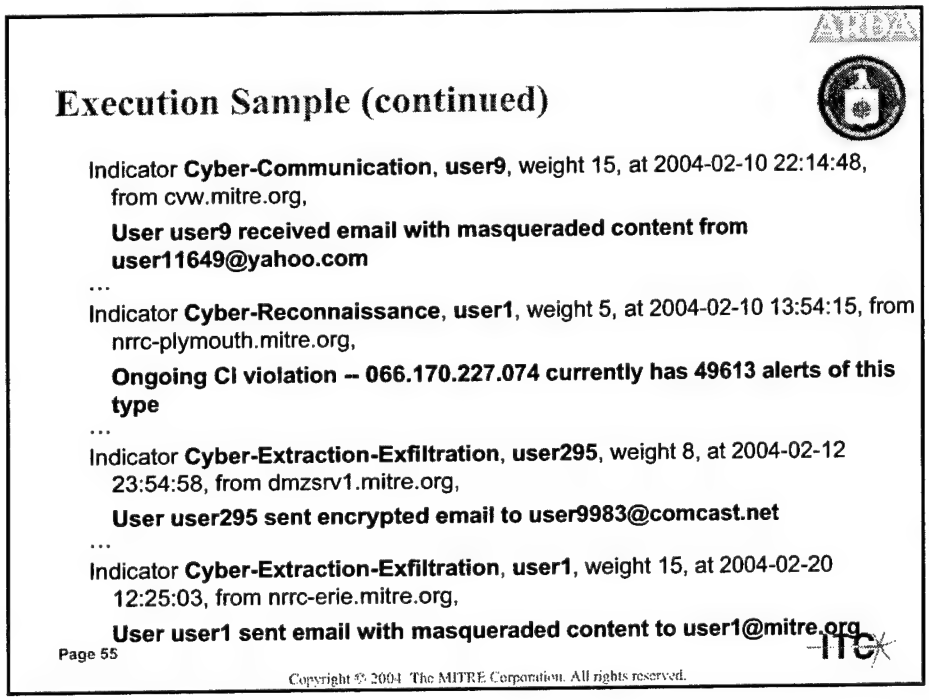
Execution Sample (continued)
Indicator Cyber-Communication, user9, weight 15, at 2004-02-10 22:14:48, from cvw.mitre.org,
User user9 received email with masqueraded content from user11649@yahoo.com
...
Indicator Cyber-Reconnaissance, user1, weight 5, at 2004-02-10 13:54:15, from nrrc-plymouth.mitre.org,
Ongoing CI violation -- 066.170.227.074 currently has 49613 alerts of this type
...
Indicator Cyber-Extraction-Exfiltration, user295, weight 8, at 2004-02-12 23:54:58, from dmzsrv1.mitre.org,
User user295 sent encrypted email to user9983@comcast.net
...
Indicator Cyber-Extraction-Exfiltration, user1, weight 15, at 2004-02-20 12:25:03, from nrrc-erie.mitre.org,
User user1 sent email with masqueraded content to user1@mitre.org
Page 55
ITC
Copyright © 2004 The MITRE Corporation. All rights reserved.

Experimental Results

User	Weight	Breadth	Watch	Alert	Categories
user295	304	5	Y	Y	Reconnaissance (75.8%), Cyber Access (13.2%), Communication (5.2%), Physical Access (3.2%), Extraction-Exfiltration (2.6%)
user8859	252	1	N	N	Reconnaissance
user1	75	3	Y	Y	Cyber Access (60%), Reconnaissance (20%), Communication (20%)
user301	70	2	Y	N	Extraction-Exfiltration (71.4%), Cyber Access (28.6%)
user2649	70	2	Y	N	Cyber Access (71.4%), Extraction-Exfiltration (28.6%)
user322	50	1	N	N	Cyber Access
user2644	40	1	N	N	Cyber Access
user2304	30	1	N	N	Reconnaissance
user2645	25	1	N	N	Cyber Access
user2647	25	1	N	N	Cyber Access
user9	15	1	N	N	Communication
user1265	10	1			
user7448	10	1			
user215	5	1			
user324	5	1			
user268	2	1			

DATA REDUCTION

- **7.4 M records examined for 75 users**
- **259 indicators for 24 users**

Page 56

Experimental Results

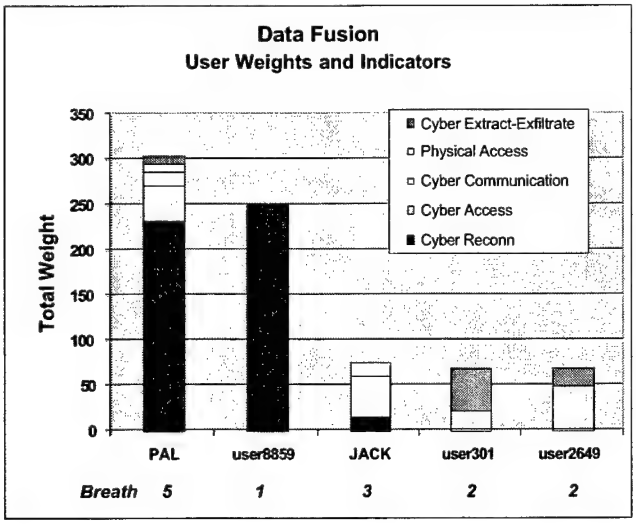

Page 57

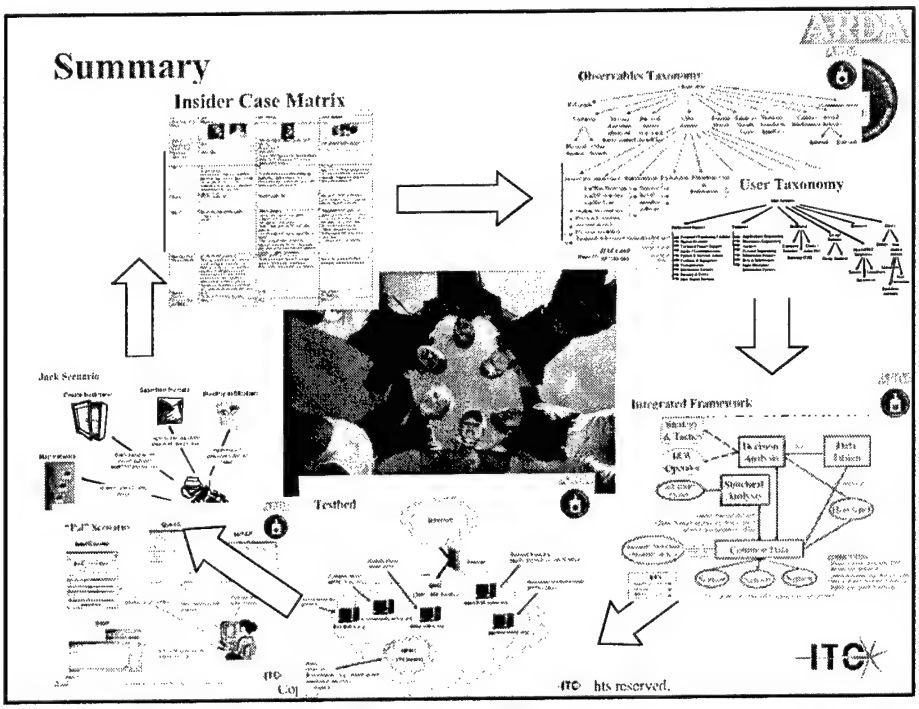

Performance Evaluation Metrics

- Developed evaluation methods/metrics
- Timeliness, e.g., time from defection to detection
  - years, months, weeks, minutes
- Accuracy
  - Precision = # correctly detected insiders / # reported
  - Recall = # reported insiders / total # actual insiders
  - False positives = 1-precision
  - False negatives = total # actual insiders - # correctly detected

Page 59

ITC

Caveat on Experiment results

- Limited users (75)
- Limited # of algorithms attempted
- Insider models motivated by actual insider behavior but require community vetting (e.g., Jack would already know network topology so would not scan)
- Differences from IC networks (e.g., no guards, open network)
- Relatively few hosts (18/400) instrumented
- Data set collected over several months, insiders known to operate over years

Page 60

Example Calculation of Fusion Performance

Identified On Watchlist \ *Actual*	MI	Not MI
MI	2 PAL, Jack	2 User301, User8889
Not MI	1 Tides-Admin	70

False Positives: 2 *False Positive Rate = 1/72 = 0.03*

False Negatives: 1 *False Negative Rate =1/3 = 0.33*

Page 61

Performance: Accuracy

	StealthWatch *			SAG		Data Fusion
	Scan	Data	Connect	PAL	Jack	
Precision	1/6=17%	1/4=25%	0%	2/8=25%	1/4=25%	2/4=50%
Recall	1/3=33%	1/3=33%	0%	2/2=100%	1/1=100%	2/3=66%
False Positive	5 (0.28)	3 (0.18)	0	6 (0.09)	3 (0.04)	2 (0.03)
False Negative	2 (0.67)	2 (0.67)	3 (1.0)	0 (0)	0 (0)	1 (0.33)

StealthWatch error rates computed on host IPs in network; SAG and Data Fusion computed on users.

Across approaches, correctly identified 3 out of 3 insiders

Page 62

ITC

Performance: Timeliness

StealthWatch
SAG
Data Fusion

Watch list
Alert

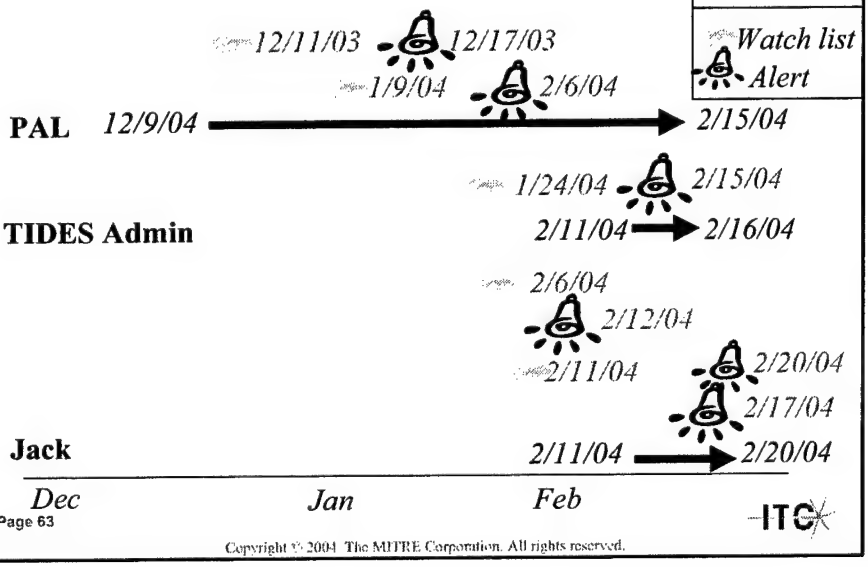

Page 63

ITC

Findings

- MI effectively represented by historical case and real experts (analyst (Penny), app admin (Laurie), system admin (Billy))
- Multisource, multilevel logging supports MI behavioral trace
  - Careful sensor deployment important (e.g., network monitoring and application monitoring)
  - Not all logs provide user or host attribution
- Cumulative non-temporal (data fusion) and temporally dependent structural analysis model both worked well
  - Both approaches insensitive to weights
  - Require expertise on alert sources and systems
- Automated broad classification of users feasible (e.g., end user versus system administrator)

Page 64

ITC

Limitations

- Unclassified, non IC network
  - Open network, no guards, no machine lock down
  - IC has swipe in/out, common networking timing
- Relatively few hosts (18/400) instrumented
- Data set collected over several months, insiders known to operate over years
- Representative insider threats addressed ... many additional models required (e.g., slow Jack attack, Ana Belen Montes)
- Common applications addressed (e.g., mail, news, web), but additional services would have enabled richer scenarios (e.g., instant messaging, database access) and analyses (e.g., social network analysis).

Page 65

ITC

Lessons Learned

- Methodology
  - Do scenarios up front.
- Modify classical sensors
- Clear where to place sensors for outsider, not obvious where you place sensors for insider
- Important not only to detect MI but also provide evidence for investigation
- Underestimated complexity of sensor selection and log analysis leaving insufficient time for experimentation
- Workshop format efficient and enjoyable: few but key physical meetings, work in between meetings, weekly telecon, common shared data

Page 66

ITC

Future Insider Knowledge and Focus

STEALTH	*Occurred*	*Not yet Occurred*
Detectable	*Robert Philip Hanssen*	*MI who attacks the network*
Hard to Detect	*Ana Belen Montes*	*Non-cyber component*
Not yet Detectable		

Page 67

ITC

Access to Insider Threat Data Set

- Corpus created including policy, scenario creation, instrumenting/capturing, archiving, anonymization, and database indexing network, application, and physical access logs,
- 11+M records
- Based on Infosec/legal/HR review, availability beyond workshop requires addressing:
  - Security Vulnerabilities
    - IP and possibly machine anonymization
    - We believe no vulnerability data (e.g., no nessus scans)
    - Recovery of network topology
  - Privacy
    - User IDs from URLs
    - Machine names from URLs
  - Approval from host owners and users for new use

Page 69

SpyCatcher Motivation

- Community-wide, corpus and metrics-based evaluation has resulted in rapid advances in areas including
  - Speech, e.g., ATIS
  - Machine translation
  - Information retrieval, e.g., TREC
  - Information Extraction, e.g., MUC
- With the development of MITRE's DMZ corpus, we have an opportunity to inject simulated (historical or projected) insider behavior and test
- Systems can then be run against this data, results compared, and error analysis performed.

Page 70

SpyCatcher Challenge Process

Select Insider Case

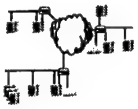

Simulate on DMZ (or OSIS)

Participants Sign Data NDA

Data Sent to Participants

Participants Submit Runs

ITC

Page 71

Detailed Plan

- Each year select an increasingly challenging insider threat (e.g., start with FBI's Robert Hanssen then move on to DIA's Ana Belen Montes)
- Annually publish the kinds of logs and indicators that will be present in the data set (e.g., no financial records but system, network, and application logs). Do this during summer workshops.
- Simulate newly modeled threat in context of real network (e.g., DMZ)
- Groups sign non-disclosure of data agreement
- Research groups are given logs and asked to produce indicators and warnings (e.g., list of suspected accounts, likely targets) in short time period (e.g., days).
- Separate categories for manual/semi-automated and automated approaches (help discover new indications and warning methods).
- Award cash prizes for the best (fastest and most accurate) and most innovative systems (e.g., $5 or $10K, plus bragging rights) as judged by panel of experts
- Share results at open workshop so community learns

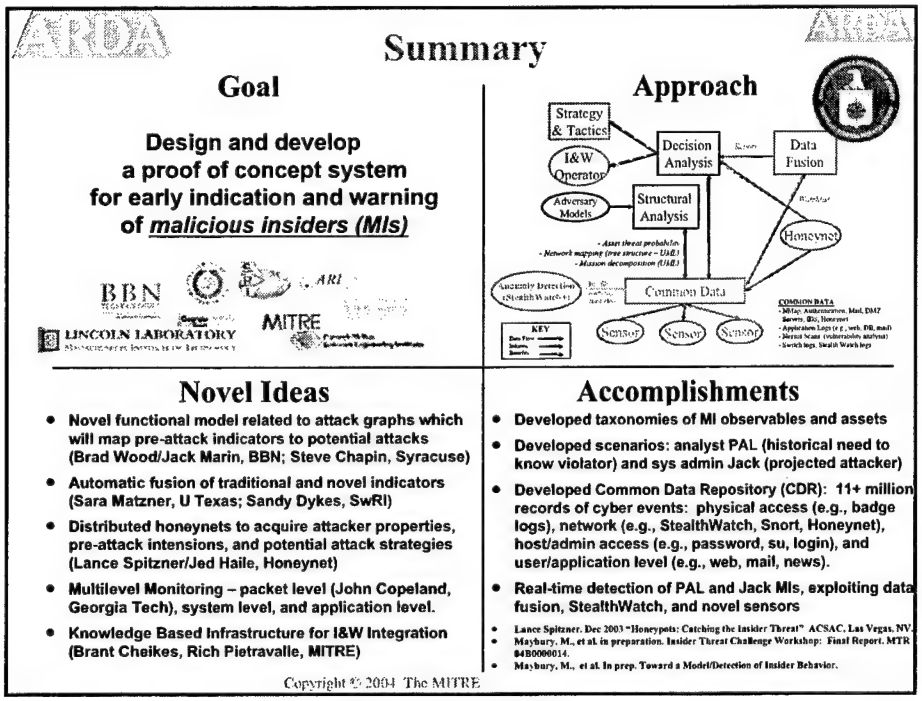
Summary
Goal
Design and develop
a proof of concept system
for early indication and warning
of malicious insiders (MIs)
BBN
ARI
LINCOLN LABORATORY
MITRE
Approach
Strategy & Tactics
I&W Operator
Decision Analysis
Data Fusion
Adversary Models
Structural Analysis
Honeynet
Common Data
Sensor
Sensor
Sensor
KEY
Novel Ideas
Novel functional model related to attack graphs which will map pre-attack indicators to potential attacks (Brad Wood/Jack Marin, BBN; Steve Chapin, Syracuse)
Automatic fusion of traditional and novel indicators (Sara Matzner, U Texas; Sandy Dykes, SwRI)
Distributed honeynets to acquire attacker properties, pre-attack intensions, and potential attack strategies (Lance Spitzner/Jed Haile, Honeynet)
Multilevel Monitoring – packet level (John Copeland, Georgia Tech), system level, and application level.
Knowledge Based Infrastructure for I&W Integration (Brant Cheikes, Rich Pietravalle, MITRE)
Accomplishments
Developed taxonomies of MI observables and assets
Developed scenarios: analyst PAL (historical need to know violator) and sys admin Jack (projected attacker)
Developed Common Data Repository (CDR): 11+ million records of cyber events: physical access (e.g., badge logs), network (e.g., StealthWatch, Snort, Honeynet), host/admin access (e.g., password, su, login), and user/application level (e.g., web, mail, news).
Real-time detection of PAL and Jack MIs, exploiting data fusion, StealthWatch, and novel sensors
Lance Spitzner. Dec 2003 "Honeypots: Catching the Insider Threat" ACSAC, Las Vegas, NV.
Maybury, M., et al. in preparation. Insider Threat Challenge Workshop: Final Report. MTR 04B0000014.
Maybury, M., et al. In prep. Toward a Model/Detection of Insider Behavior.
Copyright © 2004 The MITRE

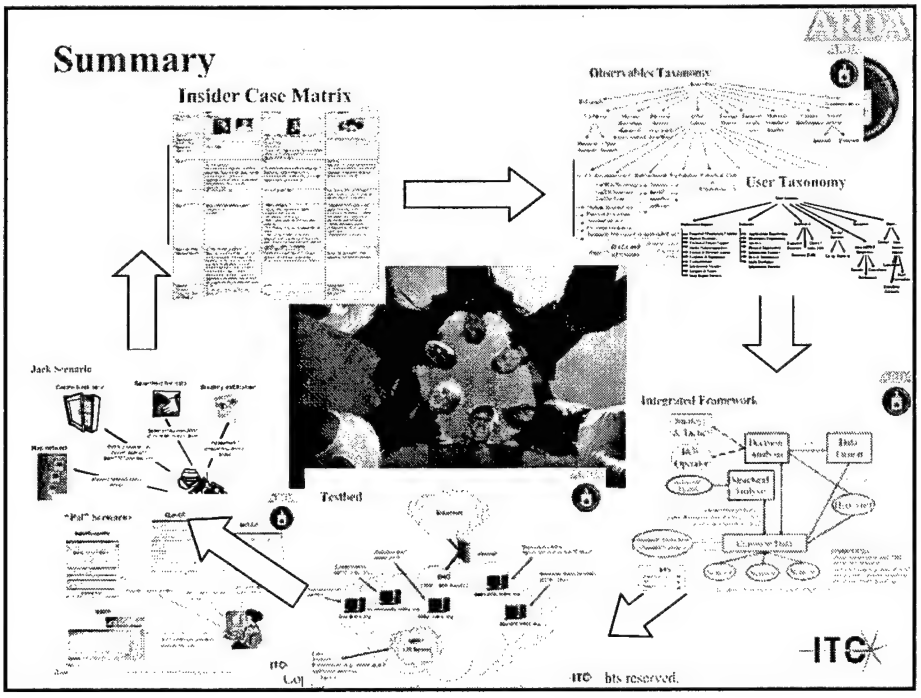
Summary
Insider Case Matrix
Observables Taxonomy
User Taxonomy
Jack Scenario
Testbed
Integrated Framework
ITC

Accomplishments

- Digital Library for Research in Insider Threat
  - Affidavits, case studies, case analysis, observable/asset taxonomy
  - Briefed Insider Modeling at ARDA's Advanced Countermeasures for Insider Threat (ACIT) Kickoff
- Experimentation Data Set and data integration environment
  - Developed Common Data Repository (CDR): 11+ million records of cyber events: physical access (e.g., badge logs), network (e.g., StealthWatch, Snort, Honeynet), host/admin access (e.g., password, su, login), and user/application level (e.g., web, mail, news).
  - Toolset to support anonymization and filtering
- Developed scenarios: analyst PAL (historical need to know violator – EO 12968) and sys admin Jack (projected attacker)
- Proof of Concept Tech Approaches for Detections
- Test Cases for Evaluation

Publications

- Lance Spitzner. Dec 2003 "Honeypots: Catching the Insider Threat" ACSAC, Las Vegas, NV.
- Maybury, M., Sebring, J., Chase, P., Chiekes, B., Pietravalle, R., Costa, M., Brackney, D., Lehtola, P., Matzner, S., Hetherington , T., Marin, J., Wood, B., Longstaff, T., Spitzner, L., Haile, J. L, Cunningham, R., Copeland, J., and Lewandowski, S. In preparation. Toward a Model of and Detection of Insider Behavior.
- Maybury, M., Chase, P., Sebring, J., Cheikes, B., Pietravalle, R., Costa, M., Matzner, S., Hetherington, T., Longstaff, T., Wood, B., Marin, J., Spitzner, L., Haile, J., Lewandowski, S., Cunningham, R., and Copeland, J. in preparation. Insider Threat Challenge Workshop: Final Report. MITRE Technical Report, MTR.

APPENDIX G

Presentation: Intelink Factoids

Intelink Factoids

Pete Jobusch, CTO
Information Assurance Directorate
Intelink Management Office
peterj@intelink.gov

Topics

- What is Intelink?
- History
- Statutory and Policy Environment
- Information Space Issues

What is Intelink?

- Intelink is NOT a Network
- Intelink is NOT a Service
- Intelink is a Corroboration

History

- Mosaic
- Intelink Services Management Center
- Intelink Management Office

Statutory and Policy Environment

- Title 50
- DCI Directives
- Community Policies

Information Space Issues

- Utility not derived from "dancing pigs"
- Not all Information equally Sensitive
- Secure COI
- Community PKI
- Other Protections

Questions?

APPENDIX H

Presentation: Glass Box Analysis Project

UNCLASSIFIED

Glass Box Analysis Project: Overview for Insider Threat Workshop

Dr. Frank L. Greitzer
Battelle, Pacific Northwest Division (PNWD)

PNWD Glass Box Project Manager: Paula Cowley

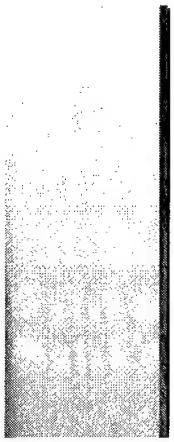

Glass Box

March 4, 2004

NIMD ARDA

UNCLASSIFIED

UNCLASSIFIED

Glass Box as a NIMD Hub

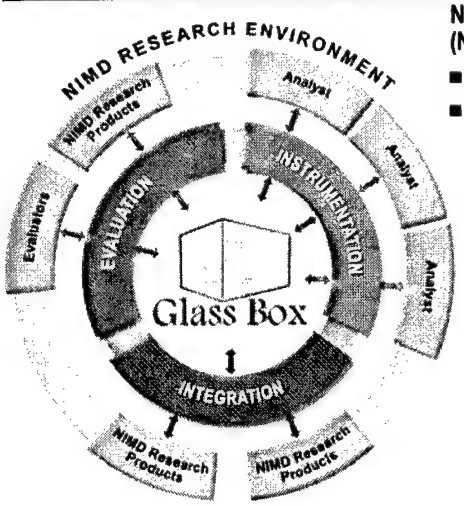

Novel Intelligence from Massive Data (NIMD)

- ARDA Research Program
- ARDA Thrust Manager: Lucy Nowell

Glass Box Objectives:

- Meaningful data for NIMD research community
- Utility software for examining data
- Facilitate new tool development by NIMD researchers

UNCLASSIFIED

UNCLASSIFIED

Goal: Capture the Analytic Process

- Collect data from signed-up analysts
  - Perform analytic taskings provided by analysts
  - Collect "ground truth" data of what analysts actually did
  - Capture cognitive thought processes
  - Capture queries and documents read
  - Capture reports generated
  - Capture "Over The Shoulder" view of activities
- Distribute Database, Filestore, and Access Tools to NIMD Researchers
- Enable reconstruction and visualization of analytic process.

UNCLASSIFIED

UNCLASSIFIED

Glass Box Instrumentation Software Automatically Captures Workstation Events

- Web Browser activity
  - URLs and contents of all pages visited
  - Images displayed on Web pages
  - Queries submitted by analyst to search engine
  - Results of query
- Application records (e.g., MS Word, PowerPoint, Excel)
  - Periodic snapshots of documents
  - Various actions such as "Find..."
- Window events (active window, location on screen, how long open, ...)
- File/Save events
- Copy/Paste events
- Keyboard and mouse data
- Etc.

Analyst-initiated annotations are also captured.

UNCLASSIFIED

UNCLASSIFIED

Glass Box Software Also Provides Review Functions for NIMD Researchers

- Tabular Review shows time-stamped activities
- "Replay" Events in 'Deja View' mode
- "Over-the-Shoulder" View mode (Camtasia)

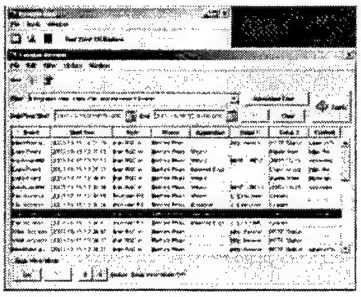

Data enable visualization of analysis process.

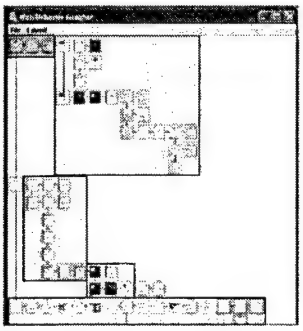

PARC's Web Behavior Grapher

UNCLASSIFIED

UNCLASSIFIED

Glass Box Instrumentation Can Be Used To Support Insider Threat R&D

Possible Applications

- Capture sample data to support event characterization and vulnerabilities R&D
- Create/capture sample data using insider threat "simulations"
- Use Glass Box API to integrate/test proposed sensors
- ...

UNCLASSIFIED

APPENDIX I

Presentation: Interacting with Information: Novel Intelligence from Massive Data (NIMD)

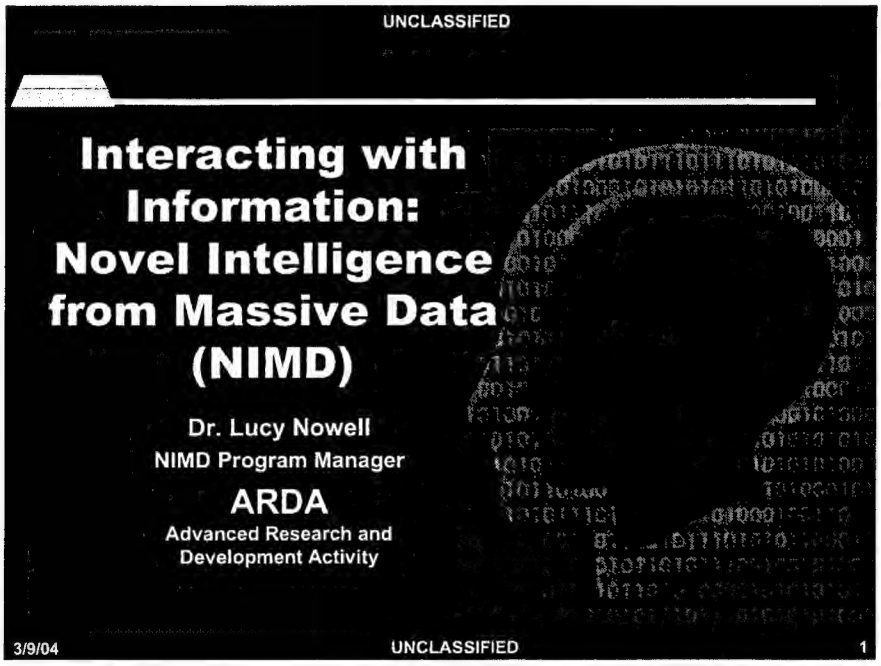

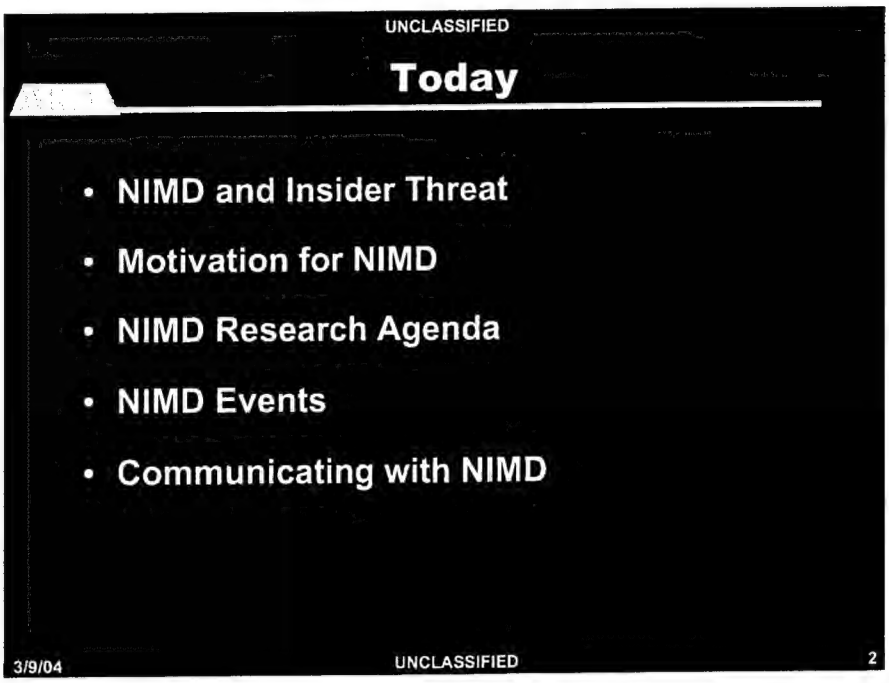
UNCLASSIFIED
Today
• NIMD and Insider Threat
• Motivation for NIMD
• NIMD Research Agenda
• NIMD Events
• Communicating with NIMD
3/9/04
UNCLASSIFIED
2

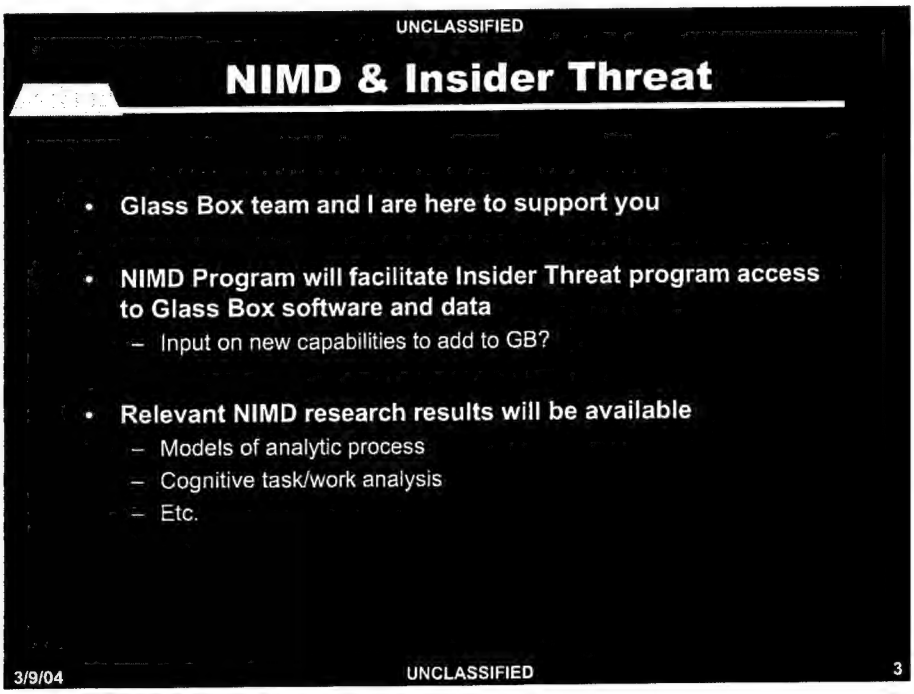
UNCLASSIFIED
NIMD & Insider Threat
• Glass Box team and I are here to support you
• NIMD Program will facilitate Insider Threat program access to Glass Box software and data
– Input on new capabilities to add to GB?
• Relevant NIMD research results will be available
– Models of analytic process
– Cognitive task/work analysis
– Etc.
3/9/04
UNCLASSIFIED
3

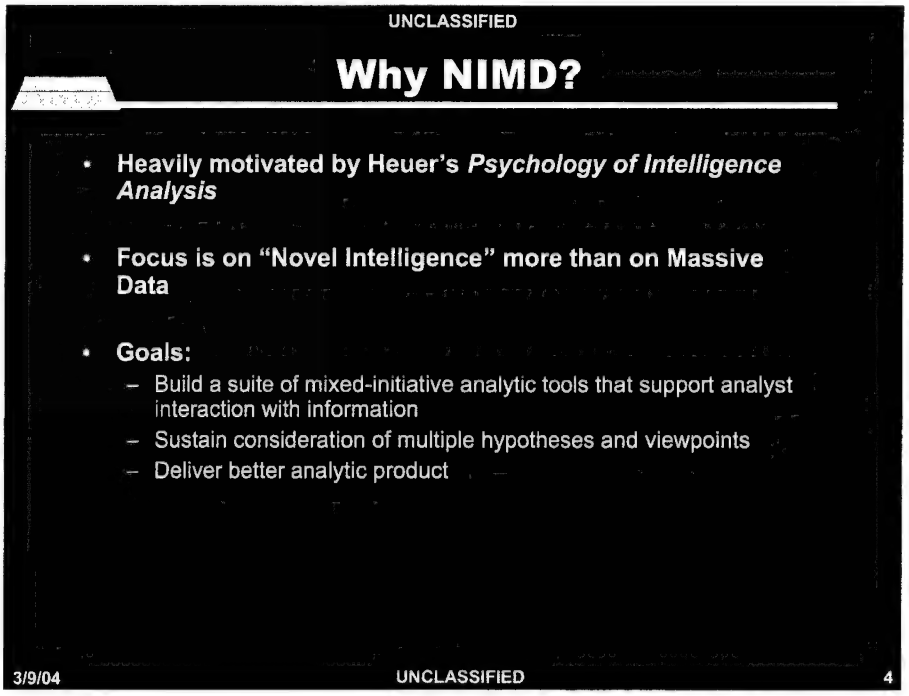
UNCLASSIFIED
Why NIMD?
• Heavily motivated by Heuer's *Psychology of Intelligence Analysis*
• Focus is on "Novel Intelligence" more than on Massive Data
• Goals:
– Build a suite of mixed-initiative analytic tools that support analyst interaction with information
– Sustain consideration of multiple hypotheses and viewpoints
– Deliver better analytic product
3/9/04
UNCLASSIFIED
4

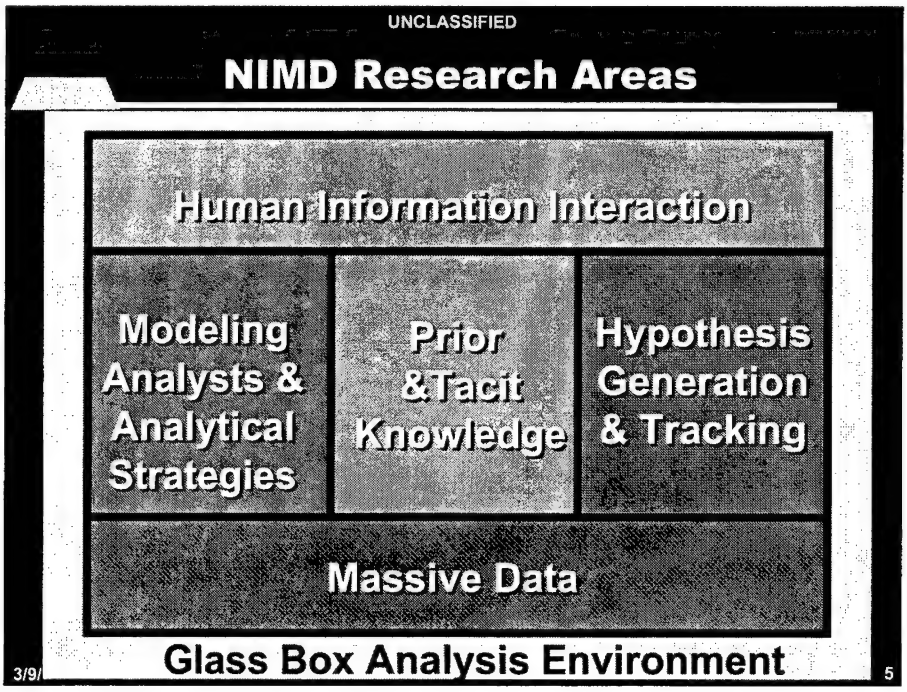
UNCLASSIFIED
NIMD Research Areas
Human Information Interaction
Modeling Analysts & Analytical Strategies
Prior &Tacit Knowledge
Hypothesis Generation & Tracking
Massive Data
Glass Box Analysis Environment
3/9/
5

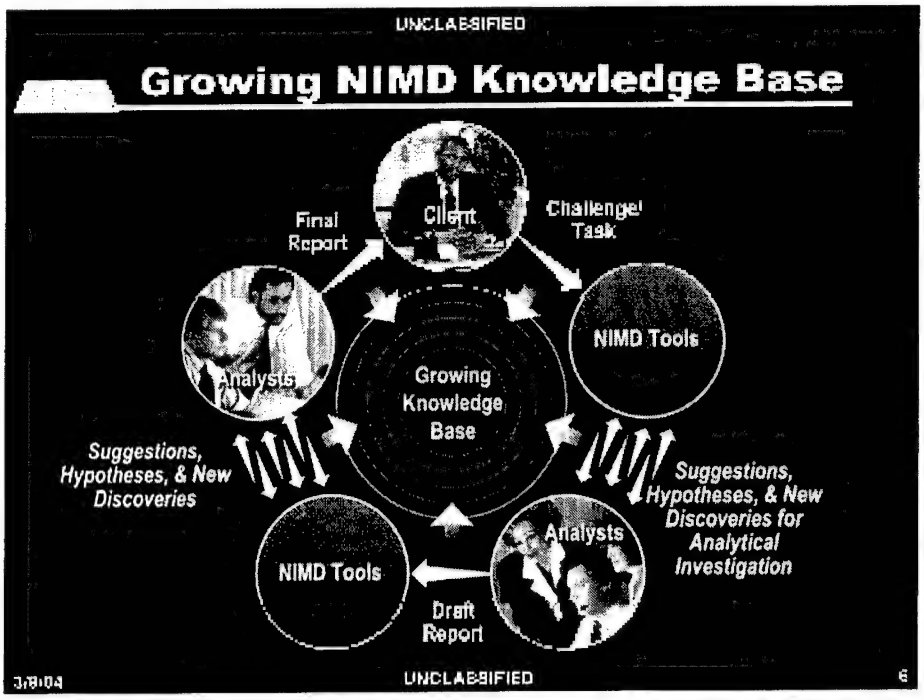
UNCLASSIFIED
Growing NIMD Knowledge Base
Client
Final Report
Challenge/ Task
Analysts
Growing Knowledge Base
NIMD Tools
Suggestions, Hypotheses, & New Discoveries
Suggestions, Hypotheses, & New Discoveries for Analytical Investigation
NIMD Tools
Analysts
Draft Report
3/9/04
UNCLASSIFIED
6

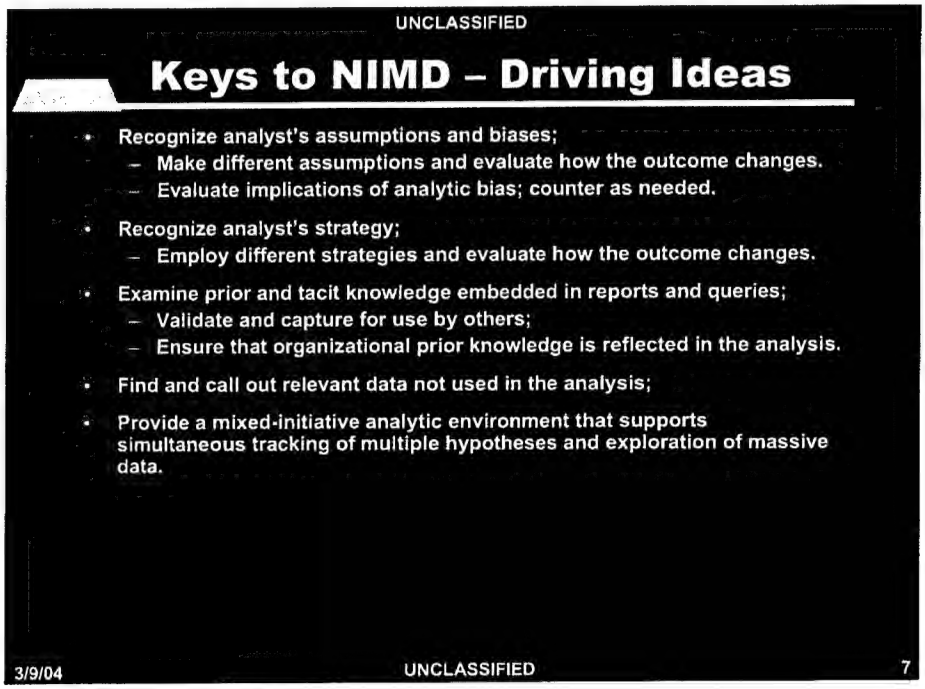
UNCLASSIFIED
Keys to NIMD – Driving Ideas
• Recognize analyst's assumptions and biases;
– Make different assumptions and evaluate how the outcome changes.
– Evaluate implications of analytic bias; counter as needed.
• Recognize analyst's strategy;
– Employ different strategies and evaluate how the outcome changes.
• Examine prior and tacit knowledge embedded in reports and queries;
– Validate and capture for use by others;
– Ensure that organizational prior knowledge is reflected in the analysis.
• Find and call out relevant data not used in the analysis;
• Provide a mixed-initiative analytic environment that supports simultaneous tracking of multiple hypotheses and exploration of massive data.
3/9/04
UNCLASSIFIED
7

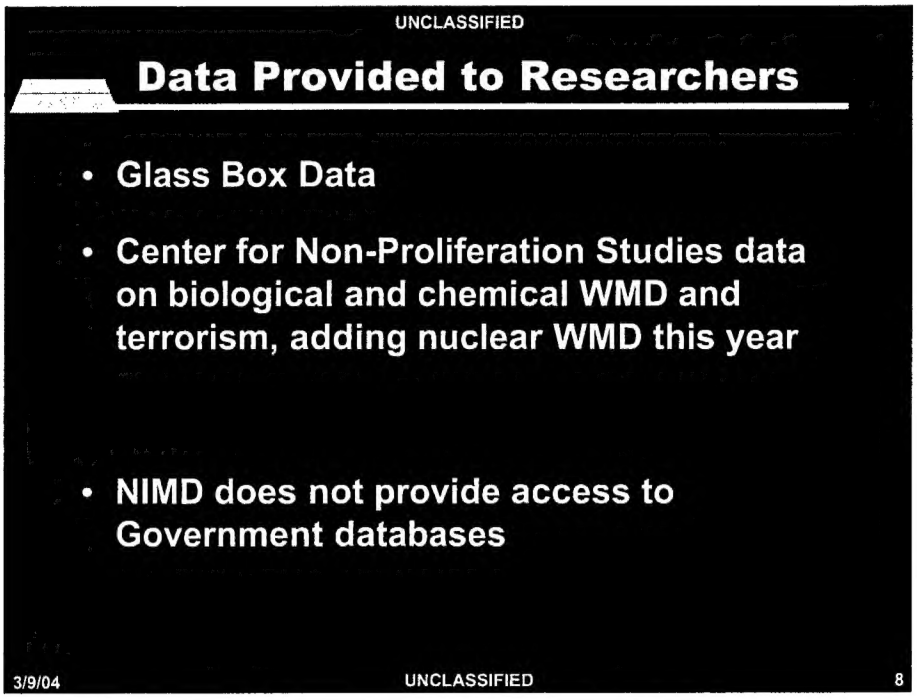
UNCLASSIFIED
Data Provided to Researchers
• Glass Box Data
• Center for Non-Proliferation Studies data on biological and chemical WMD and terrorism, adding nuclear WMD this year
• NIMD does not provide access to Government databases
3/9/04
UNCLASSIFIED
8

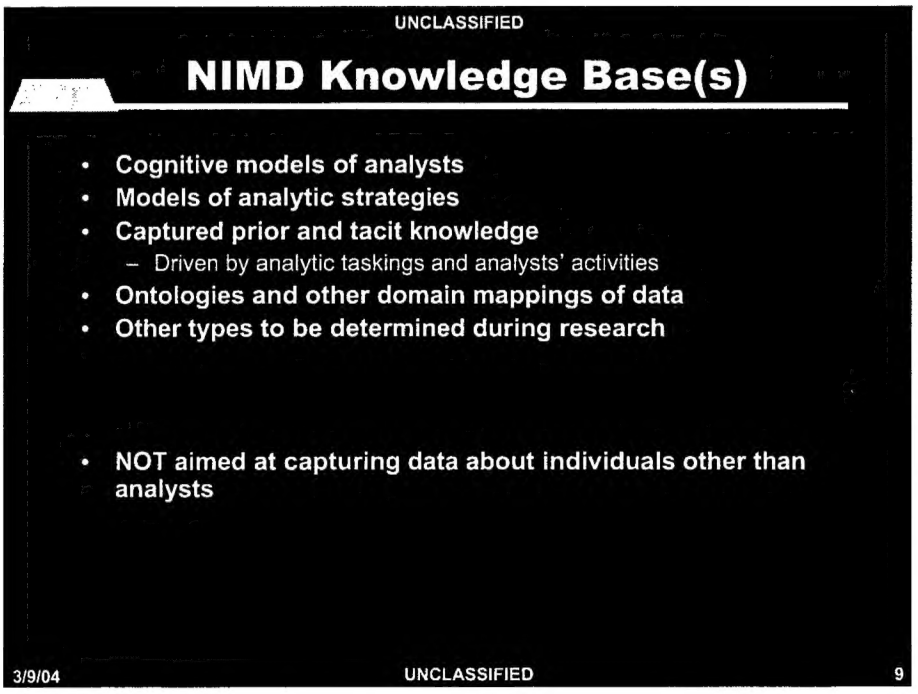
UNCLASSIFIED
NIMD Knowledge Base(s)
• Cognitive models of analysts
• Models of analytic strategies
• Captured prior and tacit knowledge
– Driven by analytic taskings and analysts' activities
• Ontologies and other domain mappings of data
• Other types to be determined during research
• NOT aimed at capturing data about individuals other than analysts
3/9/04
UNCLASSIFIED
9


UNCLASSIFIED
NIMD Events
• PI Meeting - May 25-28 in Crystal City
• Fall PI Meeting in Orlando, dates TBD
• ARDA Knowledge Representation Symposia (sponsored by NIMD)
– Three 3-day sessions
– August, October, January -- dates TBD
– Goal is to facilitate interoperability
3/9/04
UNCLASSIFIED
10

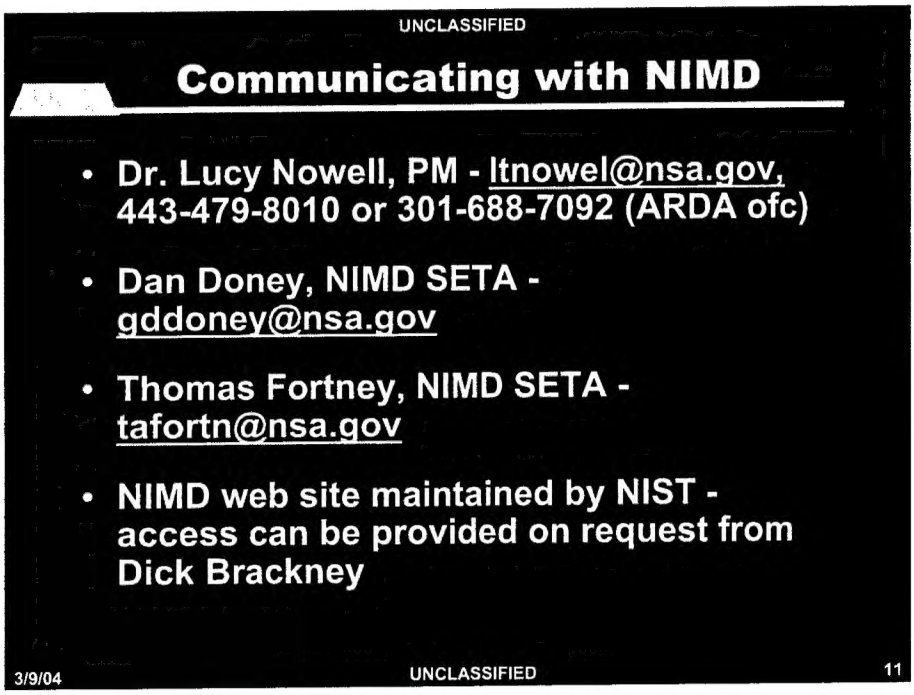
UNCLASSIFIED
Communicating with NIMD
• Dr. Lucy Nowell, PM - ltnowel@nsa.gov, 443-479-8010 or 301-688-7092 (ARDA ofc)
• Dan Doney, NIMD SETA - gddoney@nsa.gov
• Thomas Fortney, NIMD SETA - tafortn@nsa.gov
• NIMD web site maintained by NIST - access can be provided on request from Dick Brackney
3/9/04
UNCLASSIFIED
11

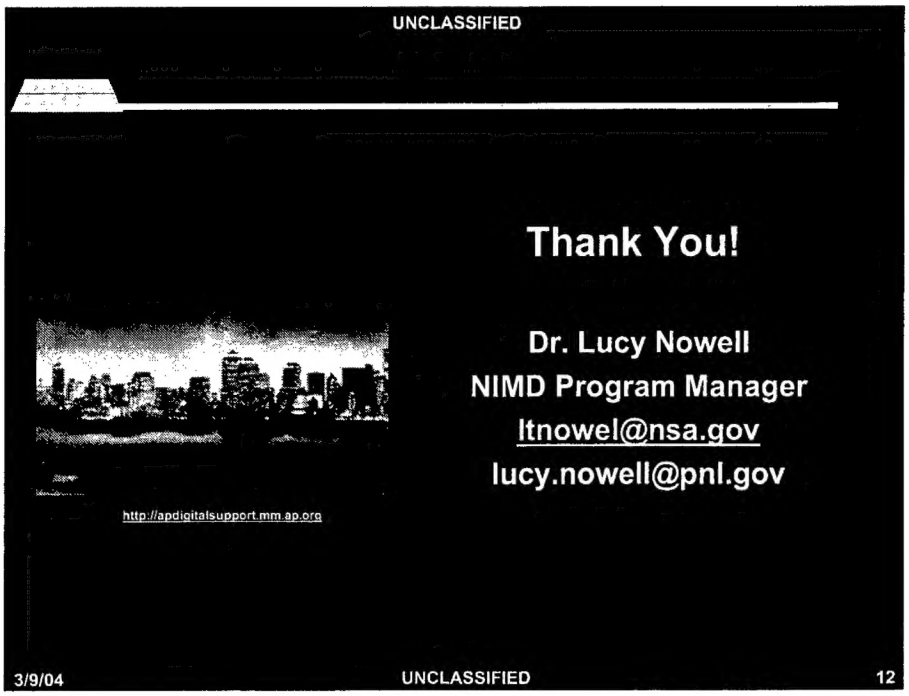
UNCLASSIFIED
Thank You!
Dr. Lucy Nowell
NIMD Program Manager
ltnowel@nsa.gov
lucy.nowell@pnl.gov
http://apdigitalsupport.mm.ap.org
3/9/04
UNCLASSIFIED
12

Bibliography

Anderson, Robert H., Thomas Bozek, Tom Longstaff, Wayne Meitzler, Michael Skroch, and Ken Van Wyk. *Research on Mitigating the Insider Threat to Information Systems-#2*: Proceedings of a Workshop Held August, 2000. Santa Monica, CA: RAND Corporation, CF-163-DARPA, 2000.

Anderson, Robert H., Richard Brackney, and Thomas Bozek. *Advanced Network Defense Research: Proceedings of a Workshop*. Santa Monica, CA: RAND Corporation, CF-159-NSA, 2000.

Commission for Review of FBI Security Programs (Webster Commission). *A Review of FBI Security Programs*. U.S. Department of Justice. March 2002. Available at http://www.fas.org/irp/agency/doj/fbi/websterreport.html.

Department of Defense. *DOD Insider Threat Mitigation: Final Report of the Insider Threat Process Team*. April 24, 2000. Available at http://www.c3i.osd.mil/org/sio/iptreport4_26dbl.doc.

Doyle, Jon, Isaac Kohane, William Long, Howard Shrobe, and Peter Szolovits. "Event Recognition Beyond Anomaly and Signature." *Proceedings of the 2001 IEEE Workshop on Information Assurance and Security*. United States Military Academy, West Point NY, 5–6 June 2001.

Heuer, Richards J., Jr. *Psychology of Intelligence Analysis*. Washington, D.C.: CIA Center for the Study of Intelligence, 1999.

Lowenthal, Mark M. *Intelligence: From Secrets to Policy*, 2nd edition. Washington, D.C.: CQ Press, 2003.

Price, Katherine E. *Host-Based Misuse Detection and Conventional OS Audit Data Collection*. MS Thesis, Purdue University, 1997.